SPECIAL REPORT

Artificial intelligence

Your questions answered

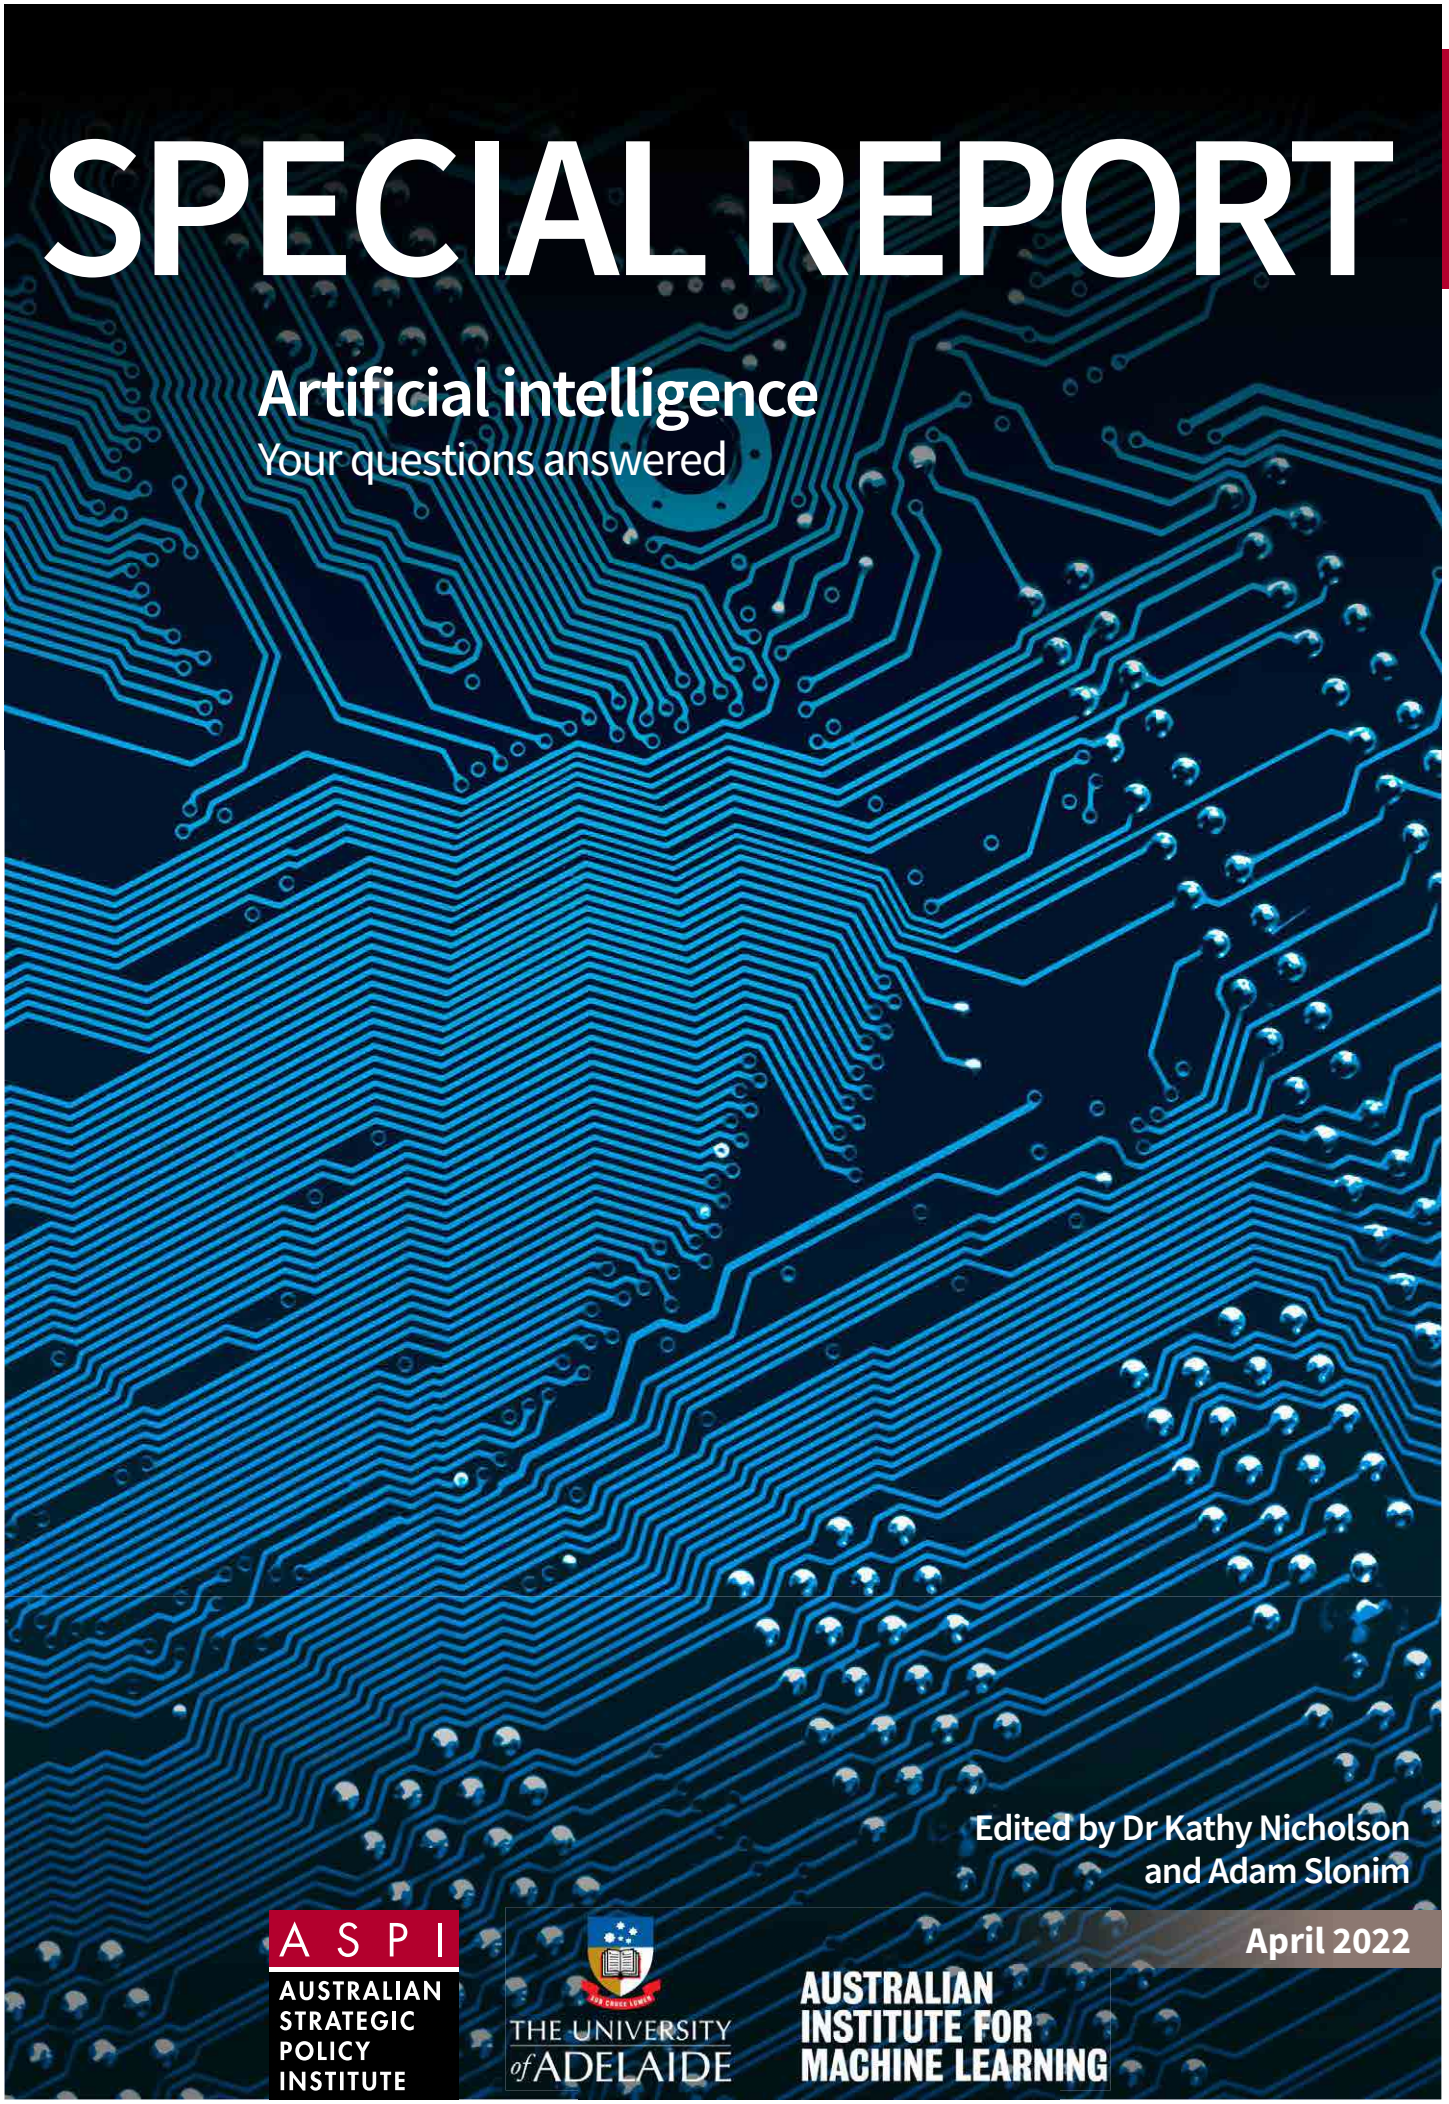

Edited by Dr Kathy Nicholson
and Adam Slonim

April 2022

ASPI
AUSTRALIAN
STRATEGIC
POLICY
INSTITUTE

THE UNIVERSITY of ADELAIDE

AUSTRALIAN
INSTITUTE FOR
MACHINE LEARNING

VISIT…

About the editors

Dr Kathy Nicholson is the Operations Manager for the Australian Institute for Machine Learning at the University of Adelaide.

Adam Slonim is the Business Development Consultant with the Australian Institute for Machine Learning at the University of Adelaide. Author of the John Curtin Institute policy report, *AI and the future of work*.

Acknowledgements

Produced by the Australian Institute for Machine Learning in conjunction with the Australian Strategic Policy Institute.

About ASPI

No specific sponsorship was received to fund production of this report.

Cover image: Electronic computer hardware, iStockphoto/Elen11.

Artificial intelligence

Your questions answered

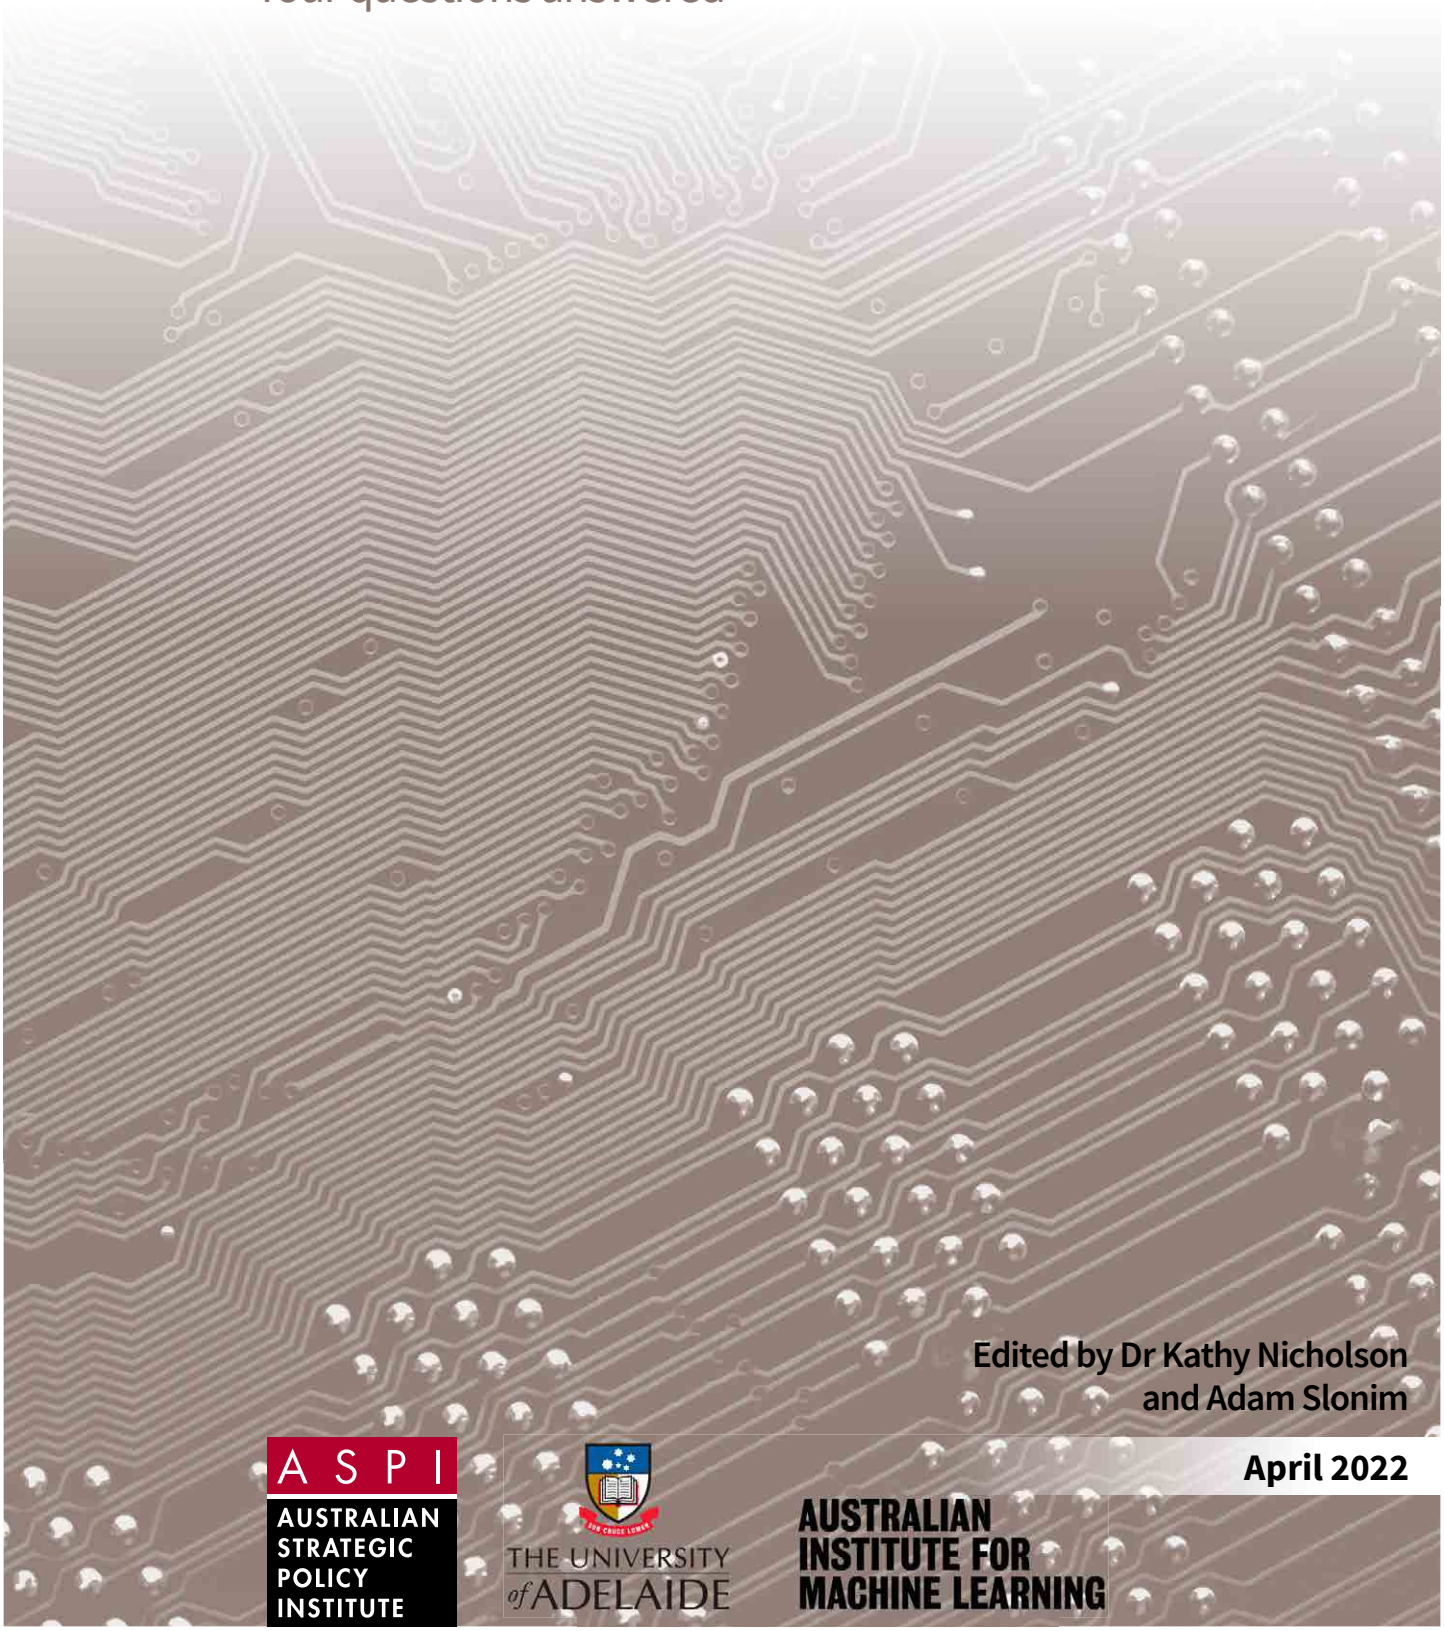

Edited by Dr Kathy Nicholson
and Adam Slonim

April 2022

ASPI
AUSTRALIAN
STRATEGIC
POLICY
INSTITUTE

THE UNIVERSITY
of ADELAIDE

AUSTRALIAN
INSTITUTE FOR
MACHINE LEARNING

First published April 2022

Published in Australia by the Australian Strategic Policy Institute

ASPI
Level 2
40 Macquarie Street
Barton ACT 2600
Australia

Tel + 61 2 6270 5100
Fax + 61 2 6273 9566
Email enquiries@aspi.org.au
www.aspi.org.au
www.aspistrategist.org.au

Facebook.com/ASPI.org

@ASPI_org

Contents

Foreword

Peter Jennings
Executive Director, Australian Strategic Policy Institute

This collection of short papers developed by the Australian Institute for Machine Learning (AIML) at the University of Adelaide and the Australian Strategic Policy Institute (ASPI) offers a refreshing primer into the world of artificial intelligence and the opportunities and risks this technology presents to Australia.

A contributor to this volume, Professor Tanya Monro, Australia's Chief Defence Scientist, writes here that 'AI could exceed $20 trillion in worth to the global economy by 2030.' That's only eight years away! AI is no longer a promising technology but is now something that's transforming all aspects of life, the global economy, social interaction, research and development and, indeed, the defence and security field.

AI's potential role in enhancing Australia's defence capabilities, strengthening alliances and deterring those who would seek to harm our interests was significantly enhanced as a result of the September 2021 announcement of the AUKUS partnership between the US, the UK and Australia. Perhaps not surprisingly, much public attention on AUKUS has focused on developing a plan 'identifying the optimal pathway to deliver at least eight nuclear-powered submarines for Australia'.

As significant as that is, there's much more to AUKUS. The grouping amounts to nothing less than a pooling of scientific, research and industrial capabilities that:

> … will enable the partners to significantly deepen cooperation on a range of emerging security and defence capabilities, which will enhance joint capability and interoperability. Initial efforts under AUKUS will focus on cyber capabilities, artificial intelligence, quantum technologies, and additional undersea capabilities.[1]

The AUKUS initiative defines a role for the Defence organisation that could be fundamentally important for AI in Australia, certainly in national security.

As Misha Schubert writes in these pages, AI is transformative, 'revolutionising our lives and workplaces'. And AI is competitive, with the global powers investing billions into research and development. Australia needs to be part of this revolution, working closely with our allies and friends, ensuring that we have thought through the ethical and social implications of AI. Australia needs some of its brightest minds working on AI, reflecting the increasing priority that government, the private sector, our allies and competitors are all putting on this area.

This AIML/ASPI report is a great starting point for individuals looking to better understand the growing role of AI in our lives. I commend the authors and look forward to the amazing AI developments to come that will, we must all hope, reshape the world for a more peaceful, stable and prosperous future.

Introduction

Artificial intelligence (AI) has arrived.

Professor Simon Lucey

Director, Australian Institute for Machine Learning, the University of Adelaide

Professor Anton van den Hengel

Director, Centre for Augmented Reasoning, the University of Adelaide

Artificial intelligence (AI) has arrived. It has become increasingly part of the fabric of our society and is embedded in so many products that we now take it for granted. AI is all around us and is becoming a key driver of economic growth. More and more services and products are having it incorporated into their design and development.

So, what exactly is AI? What does it do? How does it benefit our lives and the world? And what dangers does it present?

Many people's understanding of AI comes from sensational Hollywood movies, and of course those are fictional stories. Unfortunately, sensation and fear often take the place of grounded facts.

One of the purposes of this special report is to separate the facts of AI from fiction and myth. In so doing, we're seeking to inform the public on how AI works, how it's affecting our society and how Australia can take advantage of this remarkable technology.

The implications for Australia are significant in every sphere of human endeavour: in the jobs we do and will be doing, in the things we buy and how we buy them, in all aspects of our economy, in determining how governments can run more efficiently and effectively, and in our defence and national security.

We're also seeking to help inform the debate about AI and assist in making ethical choices when it comes to developing and introducing AI.

The team at the Australian Institute for Machine Learning (AIML) at the University of Adelaide—which includes the Australian Government-funded Centre for Augmented Reasoning—is at the forefront of developments in AI and is keen to set Australia up for success as AI seeps into every aspect of our lives.

Our aim is to inform and educate, and we hope that this report will help to build a greater understanding of what AI is and isn't, and what we can do to take maximum advantage of this technology for our future.

We thank the many contributors from science, business, government and national security who have helped shape this report.

Finally, this report could not have been possible without the efforts and support of Peter Jennings and the team at the Australian Strategic Policy Institute (ASPI). When we approached Peter with the idea for this report, his immediate enthusiasm to make it happen and help make public our contributors' knowledge of AI spurred us on.

If we've missed a question that you think is important, please let us know and we'll answer your query.

Figure 1: Established in 2018, the Australian Institute for Machine Learning (AIML) is a joint initiative of the University of Adelaide and the South Australian Government

Source: Josh Geelen/AIML.

Part 1: What's artificial intelligence?

What's artificial intelligence and what benefits does it deliver?

Professor Simon Lucey
Director, Australian Institute for Machine Learning, the University of Adelaide

AI refers to the capability of an entity or a machine to exhibit behaviours that resemble human intelligence. Machine learning allows AI to learn from data, offering unprecedented opportunities for analysis, mapping and prediction in any field for which data is available—language, vision, finance, medicine, science, space, agriculture and more.

There's a lot of excitement about AI because it can work alongside us and act as a booster for human capabilities. As a productivity tool, AI replaces some of the mundane and time-consuming aspects of work, allowing us to work smarter and more efficiently. For example, AI in medicine is creating new ways to scan images and process data, freeing up clinical staff to work on important human aspects of health, such as face-to-face care of patients.

Natural language processing is an example of an AI that's already creating benefits for many people. Based on spoken language instructions, tools such as Siri and Alexa offer hands-free capability for playing music, making phone calls and adding to a shopping list. An AI called *GPT-3* uses deep learning to produce and understand human-like text and can even write computer code.

Computer vision is an AI that gives machines the capability to recognise objects as accurately as humans can. This is creating a huge impact in fields such as autonomous vehicles (self-driving cars), defence, image searching, gaming and entertainment. Australia is one of the world's highest achievers in AI and machine learning for robotic vision.

AI is playing a huge role in discovery science, including by identifying new drugs, finding new candidates for Covid-19 vaccines and predicting protein structures.

In agriculture, AI is being used to increase crop yields by targeting water and nutrient flows according to climatic and topographic factors. It's even being used to predict crop yields, which takes a lot of the guesswork out of the production process.

AI is also playing a big role in national security, especially in protecting Australia from cyberattacks, which are occurring more frequently. It's increasingly being used to monitor, interpret and predict the actions of various state and non-state actors that seek to exploit vulnerabilities in friends and enemies alike. AI is a vital element in the arsenal of responses—both defensive and offensive—that Australia has available to respond to changing geopolitical conditions.

Critical sovereign capabilities in cybersecurity, defence, homeland security and responding to 'truth-disruptions' to electoral processes and public messaging (for example, the recent bushfires, Covid-19, 5G and so on) are increasingly dependent on AI. There's a global arms race in the development of capabilities in all those areas, and a difference of 5% in performance can have existential consequences in all cases. While Australia is unlikely to be able to compete effectively in the development of large weapons and machinery for warfare, we can and should develop

a security-focused AI sector that's globally competitive. The resulting capabilities will be essential in asymmetric warfare, whichever role Australia plays.

Many new technologies and capabilities being developed with AI will enhance the welfare of many people, including by augmenting the capabilities of differently abled individuals. For example, AI is driving the human–machine interfaces of new exoskeletons that are allowing people with impaired mobility to walk around. AI has also assisted people with disabilities to live independently. Voice-assisted AI is a major breakthrough, particularly for those who are visually impaired. It helps them communicate with others using smart devices and to describe their surroundings and circumstances.

Figure 2: In addition to its advanced AI research capability, AIML has its own engineering team to build custom AI software for clients and partners across government, industry and business

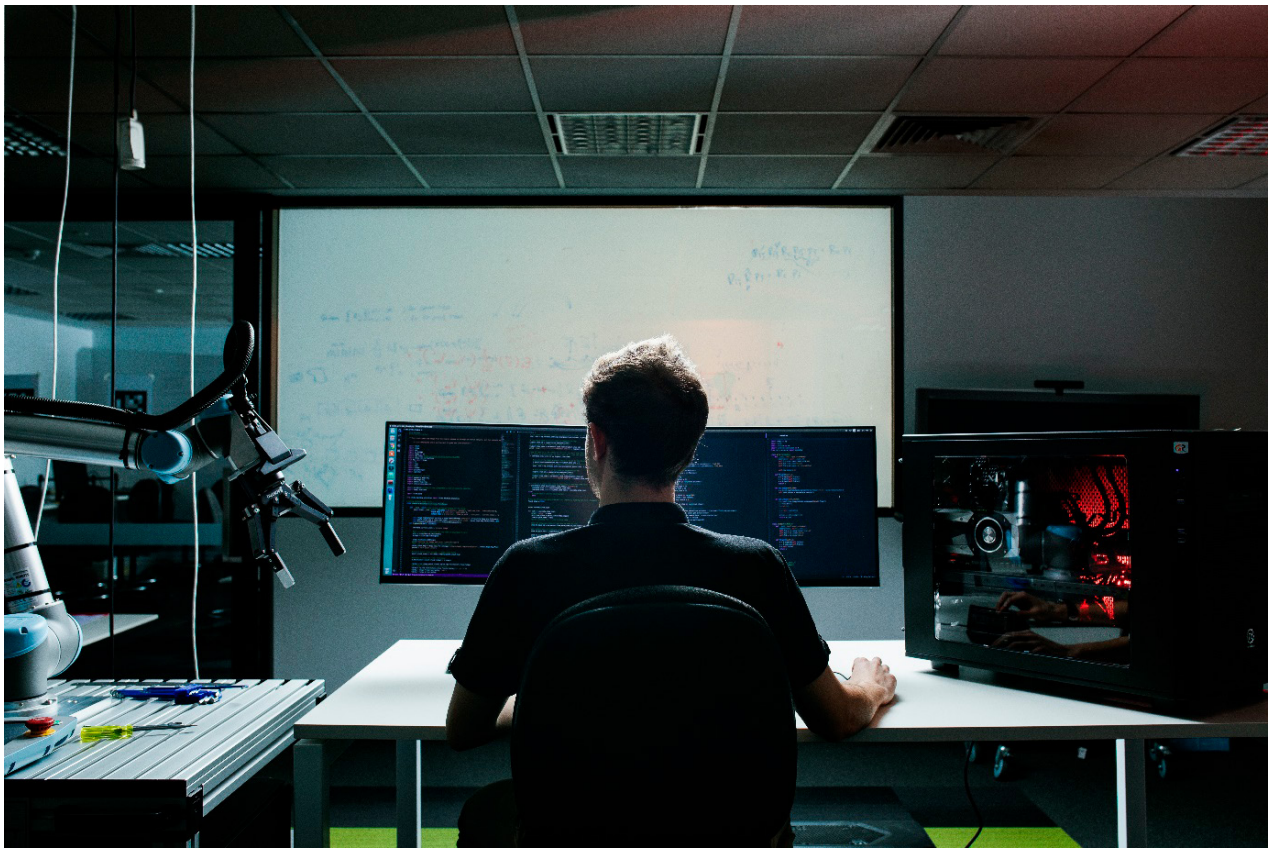

Source: Josh Geelen/AIML.

AI drives down the time taken to perform a task. It enables multi-tasking and eases the workload for existing resources. It also enables the execution of hitherto complex tasks without significant cost outlays. Put together, those two factors have radically driven down the cost of technology, enabling more technology to be available to more people than ever before. Have you used a maps program on your smartphone to voice-navigate your way to your destination along the fastest route? Everyday AI-enabled tech is benefiting people all over the planet. In a sense, AI is a great democratiser.

AI has also enabled more apps to be built to solve more problems for more people in their own communities. It's empowering more people than ever before.

AIML is involved in many projects researching new ways to develop and use AI. We work with governments, the education sector, private companies and not-for-profits in seeking ways to leverage the many benefits and advantages that AI offers.

We work with the South Australian Government and local councils using AI to investigate and deliver smarter transport networks in our cities. We work with medical providers to deliver better diagnostics, faster decision-making and better targeted drugs. And we help large organisations and bureaucracies find ways to make smarter and faster decisions that help the public with better service delivery and more effective use of our tax dollars.

The profoundly positive impact of AI is to be welcomed. While there are many myths about the risks of AI—and they aren't to be dismissed—it's ultimately human decision-making that's the key to ensuring that AI continues to deliver the many benefits it provides to society.

In summary, Australia is facing a number of challenges in which AI will deliver a number of benefits that ensure our secure, self-reliant future by:

- protecting security, society and safety
- creating jobs and prosperity
- providing better health care and doing so more cost-effectively
- maintaining our status as a global food exporter in the face of a changing climate, increasing global competition and fractious trading relationships
- rebuilding Australia as a manufacturing powerhouse.

What are the latest developments in AI?

Professor Ian Reid, FAA FTSE

Head of the School of Computer Science, the University of Adelaide
Professor, Australian Institute for Machine Learning

AI offers us unprecedented opportunities to crack old nuts with new tools—that is, to apply new technologies to solve some of the fundamental problems affecting society. In 2022, we're seeing rapid development in approaches that offer better ways to train AI through machine learning.

Foundation models are a new kind of AI trained with staggering amounts of data. They aren't stand-alone software products in themselves; rather, they form a backbone capability upon which various software tools can be built (hence the adjective 'foundation'). One example is Google's *BERT* (2018), which is a model for natural language processing and understanding trained on 3.3 billion words of written English.

The original key breakthrough in the recent surge in deep learning came from training deep network architectures from the 1980s and 1990s using what we then considered to be huge datasets. But what we understand to be 'huge' is growing quickly. *ImageNet* (2009)[2] is a database of 14 million labelled photographs that's used to test and train visual object-recognition software. *AlexNet* (2012) is a neural network architecture that has 61 million parameters. Just eight years later, another large foundational model—the text processing model *GPT-3* (2020)—is around 3,000 times that size: it has 175 billion parameters and was trained on 45 terabytes of data.

The power of foundation models lies in their flexibility to adapt to new tasks by 'finetuning' a network with a much smaller amount of data from a new domain. This is a primitive form of transfer learning—using the ability to perform one task well to transfer that ability to a new task. For example, *Dall-E* (2021) is an AI that was built on *GPT-3*'s ability to 'understand' text and combines that with the ability to generate images. It can create a picture from just an arbitrary written description of a scene. Ask it for an image of an armchair in the shape of an avocado, and that's exactly what it will give you.

The possibility of AI drawing from vast data means it has already surpassed human capability in some domains.

One example of that is computer vision. For many years, AI has been able to perform numerical or geometrical tasks that humans find frustratingly time-consuming, or simply impossible. Three-dimensional reconstruction techniques in computer vision have shown how it's possible for an AI to use a single camera moving through a scene to create a detailed and accurate map of the scene and work out exactly where it is geometrically; effectively, it's GPS from images. Humans are great at the qualitative side of this (we can easily describe our surroundings, and most of us are pretty good at finding our way back to our hotel in an unfamiliar city) but we can't generate precise, GPS-like coordinates just from our visual understanding of our environment—a bit like how we can understand some deep mathematical concepts, but a calculator will always beat us at arithmetic.

But AI is catching up to humans in the realm of semantic and qualitative understanding of the world. AI can look at one of your digital photographs and infer the geometric shape and depth of your living room. This also means that the detailed geometric maps that we build using computer vision technology can now be created with a level of understanding of what's actually in the scene, not just the 3D coordinates of the space. AI enables robots to have superhuman abilities in geometry and localisation, while now also developing some of the semantic and spatial reasoning of which humans are so capable.

This kind of spatial AI will enable autonomous systems to operate safely and effectively with humans. In the future, people will work cooperatively with machines, not be replaced by them. Advances in robotic vision research are leading to AI systems that see their surroundings by understanding images and video data in real time. Machines that can learn through spending time in real or virtual environments could help us fight fires, achieve defence goals, and inspect and maintain infrastructure in remote and hazardous environments.

But AI is of course not limited to just exploring the meaning of language and images and our relationship to the 3D world; it's also playing a vital role in responding to the pressing public-health issues of our time.

One example is an AI called *AlphaFold* (2020), which combines expertise in structural biology, physics and machine learning to predict the 3D structure of proteins based on genetic sequences. This tool is expected to revolutionise the life sciences by creating improved understanding of basic biology and revealing targeted therapies to treat disease. Researchers have already used it to make predictions about several proteins associated with SARS-CoV-2 (the virus that causes Covid-19), and scientists are right now using various AI tools to accelerate the development of new treatments and vaccines.[3]

Arguably the most significant development in AI is the speed of AI development itself. Where we once measured software production cycles in years, we're now seeing a new generation of AI capability every few months, and that tempo is showing little sign of abating.

These kinds of models can support human endeavours in language, vision, robotics and reasoning in fields including industry, resources, law, health care, the environment and education. The value of our emerging AI capability isn't in the technology itself, but in where it's applied and what it can do for the world. However, even though the capability is improving at a rapid pace, learned AI models can fail unexpectedly and harbour biases, so more research is needed to ensure that applications are ethical, explainable and accepted by society at large.

How did AI develop?

Dr Kathy Nicholson

Operations Manager, Australian Institute for Machine Learning, the University of Adelaide

The allure and potential of AI have always co-existed with humanity. AI informed the robotic Talos of Greek mythology and automata in the Middle Ages, as well as more recent figures in storytelling such as the monster of Mary Shelley's *Frankenstein*, Isaac Asimov's *I, Robot*, and *The Terminator*.

While English computer scientist and cryptanalyst Alan Turing is commonly known as the father of modern AI for his work in the 1940s and 1950s, the earliest research into thinking machines dates to the late 1830s, when Charles

Babbage designed the *analytical engine*—a concept for the first mechanical general-purpose computer. It was Ada Lovelace who wrote the first published computer program in 1843, recognising the ability of the analytical engine to solve problems of any complexity, rather than to simply crunch numbers.

A century later, in the 1940s, the first modern programmable digital computer was built and used for code breaking during World War II.

From those early beginnings, the potential for computers to solve the world's biggest problems felt tangible. But converting those ideas into reality required a series of iterative technological advances in multiple domains, including materials, computer hardware, data science and logic.

It was widely believed that the large-scale implementation of thinking machines was imminent, and that general AI was just on the horizon. Humanity needed to define mechanisms to control and regulate them. Two theories that have stood the test of time emerged:

- The Turing Test (1950) determines whether a machine can demonstrate human intelligence. To pass this simple test, a computer must engage in a conversation with a human, without being detected as a machine. To date, no computer has passed the test, although some have come close.
- Isaac Asimov's Three Laws of Robotics (1942) hold that: 1) a robot shall not harm a human, or by inaction, allow a human to come to harm; 2) a robot shall obey any instruction given to it by a human; and 3) a robot shall avoid actions or situations that cause it to come to harm itself. Variations and additions to these laws have been proposed by various researchers in recent decades.

The term 'artificial intelligence', and the birth of modern AI research, were the result of a Dartmouth College workshop in the summer of 1956. Over several weeks, around 20 mathematicians and scientists came together as a group and built consensus from what had previously been an array of divergent concepts and ideas.

AI is based on the concept that a machine can mimic the process of human thought. Two competing approaches emerged for modern AI. The first is a logic-based approach and uses rules to manipulate symbols. The second approach uses artificial neural networks that mimic how the human brain works and allows systems to be trained to solve problems.

From 1957 to 1974, AI flourished alongside rapid advances in computer hardware and technology. Computers' capacity to store information increased and they became significantly faster.

Research funding was plentiful, leading to advances in the complexity of algorithms. Early AI researchers predicted that AI would beat world-leading chess players, replace humans in the workforce and demonstrate general intelligence within 10–20 years. Early milestones included:

- 1955: the first self-learning game program
- 1961: General Motors' launch of the first robot on a production assembly line
- 1965: *ELIZA*, the first chatbot, created at the MIT Artificial Intelligence Laboratory.

But when AI failed to deliver on those early promises—primarily due to the lack of funding or computer memory and processing power—investors became disillusioned. The AI field was plunged into a series of 'AI winters', in which support for AI development all but disappeared.

Academic research continued slowly through the first AI winter (1974–1980) but was reinvigorated when a new boom in research funding in the 1980s coincided with an effort to use AI to create commercial products. Advances included:

- expert systems: the automation of a series of computer functions
- major developments in deep-learning methods
- *ALVINN*, Carnegie Mellon University's autonomous land vehicle in a neural network (1989).

The second AI winter (1987–1993) began when hardware was unable to keep up with the growing complexity of expert systems.

Emerging from that period, advances in AI, computer chips and data storage reduced the cost of using deep learning for researchers and entrepreneurs. With huge datasets, modern AI neural networks can often exceed human performance for specific tasks and can even learn from experience.

Some predictions from the early AI researchers have proved true:

- 1997: IBM's *Deep Blue* beat chess grandmaster Garry Kasparov
- 2011: IBM's *Watson* beat human *Jeopardy!* champions on live television
- Computers now reliably replace humans in tasks such as online shopping; digital personal assistants; translating language; navigation; detecting, counting, and labelling objects; facial recognition; and 3D imaging.

But other predictions, such as demonstrating human-like intelligence, still elude us. Unlike children, who can often learn from a single experience, AI needs repetition—and lots of it.

Recent advances in AI have opened up new challenges requiring input from philosophers, ethicists and legal minds to explore philosophical questions concerning data privacy, equity and regulation, as well as challenges in trusting AI systems, such as data bias and the 'black box' nature of commercial systems.

The 21st century dawned as the age that AI began to mature and deliver value to humans, at least commercially. Companies such as Amazon (founded 1994), Google (1998) and Facebook (2004) all have AI partly underpinning their phenomenal growth and success. Globally, the AI market is set to exceed US$500 billion in revenue in the next two years.

The future of AI is promising. It's clear from history that we must keep investing in R&D and work closely with communities to ensure that we can achieve the next exciting technological step.

What are the limiting factors for AI today?

Dr Paul Dalby

Business Development Advisor, Australian Institute for Machine Learning, the University of Adelaide

The factors that limit the growth and expansion of AI are very different for Australia than for the rest of the world.

Globally, there's a tidal wave of investment and activity in AI research, start-ups and existing companies. It's true to say that there's a stampede towards universities that train the talent needed to grow the AI sector and create the next generation of AI.

Australia's R&D spend per capita is much lower than for similar-sized OECD economies, and our number of patents filed per million of population, at 14 per year, is way below any other comparable economy (the OECD average is 38 patents per million).[4] This exemplifies an overall attitude to innovation that will limit Australian economic development in a Covid-normal world that requires increased self-reliance.

Australia's comparatively low R&D commitment sets us up to underperform in industries that require investment in innovation, such as AI. Characterised by processes exhibiting high levels of automation, and increasingly AI, both competitive and comparative advantages will belong to the creators of technology, not the consumers. Given the inherently opaque nature of AI, this poses a number of dangers to Australia in allowing technological control to remain with other nations, non-state actors and foreign technology companies.

Automation creates both positive and negative outcomes, and if Australia doesn't actively seek the advantages of automation—including ownership and industry optimisation—then we risk losing control of our future, and all we'll be left with is the negative effects. We've seen this happen in industries such as advertising and transport,

where Facebook, Google and Uber have employed high-value AI and software engineers in the US, and Australia has been left with fewer, lower value jobs to do the lower value work. Time is of the essence: numerous commentators, including the *Harvard Business Review*, warn that companies that wait to adopt AI might never catch up.[5]

This mismatch in investment is creating a brain drain for Australian AI talent, and many of our best and brightest in the field are being taken up by overseas, and mostly American, AI companies or AI-focused business units. The pace of development in AI makes maintaining the technical currency of staff in industry difficult. Experts can quickly lose currency and relevance if they aren't constantly keeping up with the latest developments, creating both a potential barrier to global primacy for those who don't keep up and also a barrier to new competitors for those that do. Without an ability to pay for a large internal team of fundamental AI scientists (like Amazon, Google and Facebook), the only practical mechanism for the Australian AI industry to maintain currency is an ongoing engagement with the research community—because that's where the leading-edge technological development is occurring outside of the major global technology companies. Fortunately, Australia retains some of the best AI research talent in the world, and a number of our universities are ranked in the top 10 globally in various disciplines of AI research. Australian teams recently ranked 2nd and 3rd in global competitions run by the US Defence Advanced Research Projects Agency (DARPA) and NASA, respectively. This talent is an amazing natural advantage for Australia that must be nurtured.

Despite the rhetoric about AI becoming ubiquitous, there are constraints on the speed of AI development, including the following:

- *Governance*. While the AI field is relatively new, effective organisational governance is still relatively primitive for AI. Of particular concern is data governance—who in the organisation is responsible for extracting value from a company's data, and who is protecting the interests of data contributors (such as staff and customers).
- *Fairness*. The debate about fair use of AI is still to be resolved, and new international organisations, such as the Global Partnership for Artificial Intelligence, are looking at those questions. With the rapid expansion in demand for new AI products, and the limited supply of expert practitioners, there's a strong chance that AI will be built poorly by some organisations. If developers don't understand the structure of training data, they're likely to build AI that contains inbuilt biases and errors. That will reduce trust in AI products specifically and in general.
- *Talent shortage*. There currently isn't anywhere near enough human talent to satisfy the expanding demand. This only exacerbates the issues identified above.

In Australia, compared to more advanced nations, there are additional limiting factors for AI.

The general understanding of AI, its role and its impact and benefits is relatively low in Australia (although improving quickly), so confidence in deciding to invest in AI is low. This lack of understanding has several effects that are constraining Australia's adoption of AI. There's a lack of understanding of the data that an organisation holds, the value of data to that organisation, and how to extract the enormous social and commercial gains that arise from data. This then limits the sharing and pooling of data across organisations—especially large enterprises and government—that could otherwise yield enormous value to society and the economy.

Australia's venture capital market is small and immature compared to those of similar countries. Globally, billions of dollars in venture funding is pouring into AI start-ups and scale-ups. The International Data Corporation expects spending on AI technologies to double from 2020 to 2024,[6] to reach $110 billion per annum. Australia has only a few AI specialist venture funds, and the funds they're offering are usually small—one or two orders of magnitude lower than the deals being done internationally.

In general, there's a relatively strong consensus among AI experts that Australian boards and C-suite[7] executives don't fully grasp these issues and therefore decline to invest in the technology that the rest of the world knows is yielding enormous benefit. Without investment, and the capital that's required, our AI talent is leaving Australia to follow opportunities being created overseas.

This is already costing Australia. For example, one of the greatest growth areas in AI that has immediate positive impact is health care. The challenge in most countries lies in adopting this technology into healthcare systems that are fragmented, where the teams needed to implement the capabilities aren't linked up and the technology is driven from a profit motive rather than an outcome focus. Australia has a distinct advantage in adopting AI into health care. We have a unique opportunity to leverage our coordinated national system and data, and world-class AI research expertise, and to build teams of federal and state governments, healthcare operators, researchers, businesses and capital to solve key challenges in our healthcare sector that could drive substantial improvements in cost-effectiveness and individuals' health outcomes. However, it remains difficult to attract funding for AI research. The Australian Research Council isn't willing to invest in medical research, and the National Health and Medical Research Council seems hesitant about investing in technology development. There's limited investment in basic research effort as a result. The consequence is that the Australian public isn't benefiting from potential improvements in health care, and Australian governments are missing out on potential gains in efficiency.

The limitations on greater adoption of AI in Australia are largely cultural. We haven't built a culture of investing in innovation, and we don't have a culture in our business leadership that values data and its potential through the greater adoption of AI. The good news is that there's nothing inherently limiting in terms of our geographical location or market size. We have amazing natural advantages, including well-managed national datasets and world-class research teams. A concerted effort at cultural change would enable Australia to engage more fully in this new technological revolution and reap the social, health and economic benefits as a result.

Part 2: AI and Australia's prosperity

Why should Australia feel optimistic about AI and its potential as an enabling technology?

Professor Caroline McMillen, AO FAHMS
Chief Scientist for South Australia

AI is a fundamental capability that Australia has only just started to explore. We're at the start of something really exciting, and it's difficult to even see all the boundaries of what's possible. I'm optimistic about AI. I believe it will create solutions and drive commercial opportunities for Australia.

We can shape the kinds of AI we want. For example, AI can address vital global questions by meeting sustainable development goals, developing approaches for creating and distributing renewable energy, addressing poverty, and rapidly designing vaccines targeting disease outbreaks. New technologies must work within and for societies, and for AI the broader discussions—about uses, about ethics, about regulation—have started in parallel with the technology development, which is very positive.

Diversity is one issue that needs to be addressed in AI and the core skills that surround it. Australia can improve how we place science, technology, engineering and maths (STEM) opportunities in front of talented individuals from the many parts of our society. We can change the stereotypes of what a mathematician looks like and ideas about who can become a computer scientist. We need more women in AI, more Indigenous Australians, more culturally and linguistically diverse people. Change can and does happen when we apply resolve and solutions at scale.

For teens and young adults right now, AI is theirs to shape. Gen Y and Gen Z are motivated to apply technologies to solve massive problems, such as climate change. From my experience working with these generations, we're in very capable hands.

Can't we just buy AI solutions 'off the shelf'?

Dr Paul Dalby
Business Development Advisor, Australian Institute for Machine Learning, the University of Adelaide

Investing in the development of in-house AI can seem like a daunting task. There are few people with the skills to understand what's possible, let alone build it. In the absence of investment in the development of AI, off-the-shelf AI solutions can be an alternative.

But it's important to note that buying AI isn't like buying traditional software services. The fundamental strength of AI is that AI algorithms are designed for and trained on specific sets of data. Generic AI algorithms and products are built on generic data and might not be optimisable to solve a particular problem. They may also have unknown inbuilt biases.

Figure 3: Australia needs homegrown specialist AI solutions for sectors in which we want to maintain sovereign control, such as agriculture

Source: iStockphoto, online.

AI delivers the best results when it's designed by experts who understand how to develop and adapt AI algorithms for specific datasets, can test data for biases and can then test the final product for performance and bias.

Because AI is a new technology, and we're still working out what we can expect from it, it's easy to assume that all AI is largely the same. That isn't true, and there have been some spectacular failures in the development of AI systems that then produced perverse or meaningless results. Using trusted developers to create a unique AI system for a particular scenario and obtaining third-party testing of performance standards is often a requirement for building trusted AI.

Constructing a machine-learning system that can accurately learn from often messy data to deliver accurate and unbiased predictive tools, computer vision systems or language analysis tools is both an art and a science. There's an enormous and potentially critical difference between someone who merely downloads an existing algorithm from an online library and runs it over a dataset and a specialist who deeply understands the structure of data, statistics, maths and coding solutions and has a creative ability to find novel solutions to a problem. That's why the salaries of the best AI engineers are so high.

The difference in performance and trustworthiness between off-the-shelf solutions and highly tuned AI solutions is due to a number of factors, including data quality, but particularly the quality of the data engineers and AI developers. Australia is desperately short of them, as is the rest of the world. It takes many years of training to produce world-class engineers. Demand is already far outstripping supply, and the growth in demand is likely to continue to outpace the growth in supply for some time yet. Australia must ramp up its investment in training highly skilled technical specialists in AI and data science at the VET, undergraduate and postgraduate levels. My back-of-the-envelope calculations suggest that we currently graduate 100 or fewer PhD students in high-end

machine learning per year in Australia. I suspect about half of those get jobs overseas, leaving 50 people per year to be shared among the businesses, government agencies and start-ups that need them. In a Covid-19-afflicted world, it's harder to fill those gaps with international workers. The solution is a massive investment in this high-end specialisation and a national campaign to attract people to this career.

Australia particularly needs homegrown specialist solutions for sectors in which we want to maintain sovereign control. Those sectors include defence and national security, but may also include some of our large agricultural industries, critical minerals and health care (as examples). It would be unacceptable to most Australians to have those industries end up becoming 'Uberfied'—where the control of the sector is in the hands of foreign technology companies that collect all the data, undertake all the data analysis, and issue instructions to low-paid Australian workers on what to do next.

For defence and national security, there's a global battle for supremacy in AI. Most of the AI needed will be for cybersecurity, logistics, managing the data deluge and responding to 'truth-disruption' of electoral processes and public messaging. Australia should seek to develop and maintain advantage in these areas. It might seem surprising, but US defence agencies have visited Australia to learn about our superior capabilities in AI for security and related areas. We're eminently capable at holding our own and developing world-class systems.

There's a global arms race in the development of capabilities in all these areas, and a difference of 5% in performance can have existential consequences. While Australia is unlikely to be able to compete effectively in the development of large weapons and other large machinery for warfare, we can and should develop a security-focused AI sector that's globally competitive. For example, another middle power, Israel, has successfully developed a world-class capability in cybersecurity by investing in skills development in its army and universities, and then encouraging those staff to spin out cybersecurity companies through grants, start-up funding and contract purchasing. Australia has the capacity to be globally competitive in AI, which would have the double advantage of giving us superior capabilities in our national services and creating a new and globally competitive industry sector.

Buying off-the-shelf AI solutions might seem like a cheaper solution than developing unique systems, but it may also end up being a false economy if it doesn't work, or isn't trusted by its users. For industries with sovereign strategic value, it makes even more sense to build our own systems here in Australia. That won't happen without substantial new investment in high-end skills. With that investment, the result will be better, more trusted AI, but it will also support highly paid jobs and build capability that will sustain Australia's high standard of living in this new industrial revolution.

Where do Australia's strengths in AI currently lie?

Michael Evans
Evans AI

Australia possesses a wide variety of strengths across the AI value chain that, if appropriately supported and amplified, could create areas of competitive advantage to establish the nation as a global leader in AI.

At the foundational layer, the nation's strengths include modern infrastructure, a maturing data ecosystem and a strong academic and research sector. Our greatest strength in AI lies in the capability and technical layer, in which a world-class workforce of AI researchers is advancing state-of-the-art AI in cyber-physical systems, robotics and computer vision.

Australian AI start-ups (while limited in number), as well as AI institutions (such as CSIRO's Data61) are also strengths in AI capability. Broader environmental strengths, including a technology-accepting consumer market, comparatively stable political environment and attractive lifestyle factors, can also be leveraged to aid Australia's AI industry. Perhaps unsurprisingly, given the composition of the nation's economy, current national strengths at the application layer cut across resources, health care and agriculture. Australia also has a valued strength in its

reputation for developing trustworthy and responsible AI applications, and is currently working towards developing strengths in AI governance, notably via the development of standards.

The combination of those strengths provides a solid baseline from which Australia can mature its AI capabilities and position itself to become a global leader in the advancement, development and deployment of AI.

How can AI build a stronger Australian society?

Dr Catriona Wallace

Chief Executive Officer, Ethical AI Advisory

Currently, we know that an average Australian middle-aged person will interact with an AI touchpoint around 28 times per day (for a teenager, it will be more than 100 times a day[8]), so AI is already a key part of society. The software robots in our pockets—our smart devices—are essentially extensions of our biological brains and could be regarded as active members of society already. Most people don't realise that, but, when I'm on stage, I acknowledge the humans and the robots that are present.

The vast benefits that society will experience from AI will be centred on better health care, smarter cities, greater security and the personalisation of … well, just about everything. AI will make society stronger by providing better predictions and more rigorous decisions. It will enable innovation and drive productivity gains.

Within the next few years, AI will also be the foundational component of the majority of software that society already uses. It will be a source that powers the engine room of society—its data, its infrastructure and its communication channels.

And, as society is in a current state of crisis, AI will enable better communications, logistics and signal detection and help facilitate collective action to better predict and handle crises.

Thanks to AI, society's systems will be more capable, ubiquitous, accurate, efficient, smarter, seamless and predictive.

Another key and recent trend that will be AI-driven and have the potential to create a far more connected Australian society is the coming of the metaverse.

The metaverse is an online virtual world that incorporates augmented reality, virtual reality, AI, 3D holographic avatars, video and other means of communication. The metaverse is regarded as a hyper-real alternative world for society.

Elements of the metaverse already exist in applications such as Fortnite, Minecraft and Roblox. Players can already attend concerts, travel around the world, attend conferences and play games in this online space.

What if Australia were to play a key role in establishing the metaverse? Would Australians who perhaps don't have access to many services and experiences in 3D life then have the opportunity to be more connected and more active in society—albeit a virtual society? Those are great questions that we'll need to answer very soon. I strongly believe that if Australians can embrace, adopt and, most importantly, trust AI, then Australian society will be much stronger.

Part 3: What's happening in the rest of the world?

How can Australia build sovereign capability in AI, and why is it that important?

Professor Tanya Monro, FAA FTSE FOSA FAIP GAICD
Chief Defence Scientist, Defence Science and Technology Group

AI is already ubiquitous in our lives and will be foundational to our future prosperity. In Australia, the Department of the Prime Minister and Cabinet has identified AI and machine learning as critical technologies of national interest.[9] AI applications are wide-ranging and encompass many sectors, including agriculture, education, energy, finance, health care, manufacturing, transport and telecommunications, among many others.

As a general-purpose technology, AI applications in those domains may also have implications for defence and security. Indeed, the Australian Defence Force's ability to understand the operational environment, manoeuvre and project force will be transformed through advances in AI and human–machine partnerships,[10] employing dual-use and defence-specific AI applications.

Innovation in AI is accelerating, creating opportunities to grow the economy, solve challenging problems, accelerate discovery and enhance national security. It's estimated that AI could exceed $20 trillion in worth to the global economy by 2030.[11] This potential for transformative impact, and the competitive advantages offered by AI, are being actively pursued around the world. By December 2020, more than 30 countries and regions had published AI strategies or similar documents; the world's top universities are increasing their investment in AI education; and private investment in AI continues to grow.[12] This ongoing expansion in AI activity will bring accelerating advances across the field.

With global AI investments growing, Australia will need to remain internationally competitive in a world of expanding AI capabilities and AI-enabled operations. For defence, AI will be critical in delivering strategic objectives and maintaining a capable, agile, and potent defence force.[13] Reflecting that importance, and the need for increased self-reliance, AI is included among the new 'sovereign industry capability priorities' critical to defence. A sovereign capability in AI ensures Australia's stake in a key 21st-century industry and underpins the Defence organisation's future operational and training capabilities.

Australia has some world-leading AI capabilities in universities, research organisations and industry. To expand from that base, we need to focus on AI and broader STEM initiatives to build the specialist workforce required by Australian industry. Building human capacity in AI will be critical, as demand is already outstripping supply, and the need for AI specialist skills is expected to grow significantly.[14] For sovereign activities, including sensitive work supporting the nation's defence and security sector, the ability to draw upon an Australian AI-skilled workforce will be fundamental. Capacity in the Australian workforce will ensure that we can maintain an effective defence force and develop the future capabilities needed to meet and overcome the disruption that AI will introduce to current methods of warfare. That will require sustained investment in skills training at all levels to raise understanding of

AI by managers and to further expand AI graduate and higher degree by research programs to attract and train Australian AI specialists.

Increased research and technology development, supported by long-term and sustained investments in data, tools and critical digital infrastructure, will underpin our national AI competitiveness. Through sovereign investment in AI R&D, the benefits of AI will be fully harnessed and the potential risks mitigated. Investment should include building access to sufficient computing capacity and the fast networking needed to source those capabilities. In addition, data holdings need to be established and managed to support the development of AI applications and tools, buttressed by appropriate access and privacy protections. Issues of safety, security and bias need to be understood and actively researched to support and inform responsible AI practices and build trust in AI systems to underpin confidence in continued AI development. That will require the development of AI systems that accord with Australian values and are trusted partners across the breadth of AI applications. The ongoing development of AI standards, engaging a broad cross-section of stakeholders, will enable the widespread use of responsible AI and increase the competitiveness of Australian industry.[15]

Figure 4: With global investment in AI increasing, Australia needs to remain internationally competitive amid rapidly expanding AI capabilities

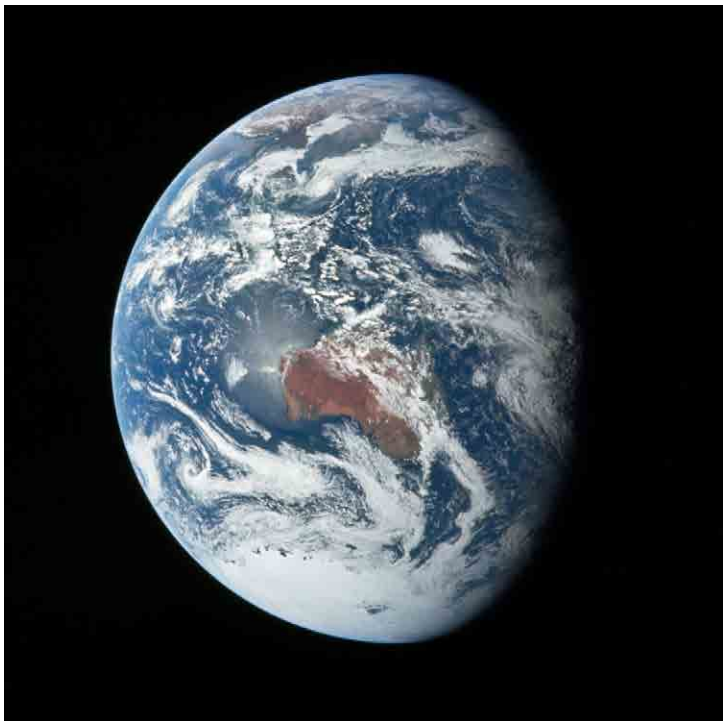

Source: NASA, online.

Through government leadership, a strong and vibrant AI ecosystem should be established to bring together partnerships with industry and academia to drive innovation, meet national priorities and expand AI exports. This will ensure strong stakeholder participation and dialogue to inform future policy considerations and connect industry and academia with new business opportunities. Efforts are already being made in that direction. In the defence sector, this includes the establishment of the Defence Artificial Intelligence Research Network (DAIRNet). The DAIRNet will work to build a research community in Australia supporting Defence, bringing together multidisciplinary teams to address AI challenges at scale. Beyond research and innovation, a sovereign capability entails the adoption and continuous upgrade of AI technologies to enhance productivity, requiring industry participation across the AI product life cycle.

Building a sovereign AI capability will enhance Australia's standing as a leading contributor to AI innovation and development, and as a credible contributor to international forums. It will ensure Australia's participation in international cooperation through substantive contributions, including in R&D, to address AI challenges, ensuring that Australia's voice is heard on key issues related to AI. Depth in Australian AI skills will underpin the successful integration of imported AI capabilities and, supported by a vibrant ecosystem sustained through investments in critical infrastructure and R&D, will attract international suppliers to use Australian industry.

What opportunities does AI present for Australia?

Professor Simon Lucey

Director, Australian Institute for Machine Learning, the University of Adelaide

Building sovereign AI capability is vital for Australia. AI will support and grow the industries our economy relies on, create new opportunities and help us sit at the global table with other high-achieving AI nations. It will help us complexify our economy. Through automation and other technological capabilities, AI will boost our nation's productivity and overcome many limitations of our relatively small population. To achieve those goals, Australia must excel in AI, and that requires significant new investment.

Australia is a huge country that relies heavily on road and rail networks to distribute food, fuel, minerals and other goods. AI and machine learning offer significant opportunities to optimise and automate transport and logistics, increasing efficiency and providing new driving technologies.

Managing the impacts of climate change is another key capability that Australia can build through developing our own AI. Smart algorithms and robots will boost our capacity to predict and control bushfires. Sophisticated algorithms will help us monitor coral reefs and oceanic conditions, manage the environment and support economies linked with tourism and aquaculture.

AI will offer us new ways to manage diminishing water resources on land and maintain agricultural outputs, including deciding what to plant and when, predicting yields and monitoring livestock.

Australia is facing a number of headwinds in maintaining our status as a global food exporter:

- Climate change is eating our productivity gains from new innovation, meaning that Australian agriculture's global competitiveness is flatlining.
- There's increasing global competition for our global markets. For example, the Grains Research and Development Corporation projects that Australian wheat might not be able to maintain its price competitiveness with grain from Eastern European countries, meaning that we would no longer sell wheat into a commodity market.
- China as a key trading partner is targeting our agricultural sector to punish Australia for perceived diplomatic offences.

'Business as usual' seems to be out of the question. We need to find new ways to value-add Australian produce to sustain our local agricultural industry sector. However, there are also many ways we can use AI to better target high-value products, reduce costs and improve quality.

Traditional models of investment in R&D in agriculture have focused on understanding biological systems and processes. By its very nature, that requires investment in research over multiple seasons to collect sufficient evidence to verify improved production capabilities. This has resulted in remarkable improvements in both the productivity and resilience of Australian agriculture. It is, however, a less than perfect model for developing technology such as AI for the industry.

The rapid pace of development in AI makes it unlike other areas of agricultural research. AI technology exists as algorithms on general-purpose hardware. Entirely new capabilities can thus be deployed on a farm in the time it takes to download new software. We have the potential for achieving world-leading capabilities to assist farming systems using AI over a relatively short period, given the right investment models.

There's an opportunity to develop 'hothouse' AI labs for agricultural industries around Australia for the rapid development of AI and machine-learning prototypes for agriculture, adopting a learning-through-doing approach.

Such hothouses would bring together world-leading researchers, machine-learning engineers and agricultural experts in a creative and fast-paced environment. The team should use a 'fail fast' mindset to quickly deliver high-tech minimum viable products, prototypes, feasibility studies and experiments in the machine-learning space that are needed by agriculture, greatly accelerating the traditional innovation cycle.

The recent Covid-19 crisis has resulted in a rethinking about the importance of local manufacturing for strategic supplies. Government, industry and communities have realised that there's a sovereign risk in outsourcing to overseas companies the manufacturing of key strategic goods—medical supplies, food, safety, transport and so on.

Australian manufacturers struggled to survive the 'Dutch disease' of the last mining boom, when the Australian dollar and wages accelerated rapidly and substantially. Whole manufacturing sectors were wiped out as a result. This was particularly the case for manufacturing, in which Australia did not have or had lost its global competitive advantage before the mining boom took off.

There's now an opportunity to rebuild a manufacturing sector that can sustain high wages, grow jobs and compete internationally. A key capability that will support this will be AI, which will:

- suppress the costs of production through improved operational efficiency and greater automation
- improve the perceived value of products manufactured with AI embedded and create barriers to competition.

There's an opportunity for the manufacturing, government, university, VET and research sectors to partner to:

- train a new generation of manufacturing workers and managers who are confident about and competent in integrating AI into their businesses and products
- co-develop new AI systems that improve operational efficiency, including optimisation software, monitoring systems and robotics
- co-develop products with AI embedded in them that will be in global demand.

The mining and space sectors are increasingly dependent on AI, as are defence capabilities, including weapons development and surveillance. It's critical that we develop our own capabilities in those areas if we're to learn the lessons of Covid-19 disruptions to global supply chains and China's more adversarial approach to its economic relationship with Australia.

What are AI superpowers doing to advance their incorporation of AI into their defence postures, and what are the national security implications for Australia?

Michael Shoebridge
Director, Defence Strategy and National Security Program, Australian Strategic Policy Institute (ASPI)

In 2022, there are two kinds of AI 'superpowers': companies and states. The most capable AI national security power will be the state that has the closest connections to and is able to take advantage of corporate AI superpowers. AI capability isn't simply transferable from one application or sector (for example, search, facial recognition or digital navigation) to another without deep understanding of the uses and purposes involved and the limitations of the available datasets and resulting machine-learning applications. AI uses in national security look compelling and potentially destabilising, from insight advantage from huge datasets to the control of autonomous systems and rapid decision-making.

Right now, the US is a potential AI superpower, and much of its technical capability is coming from US big tech (notably Amazon, Apple, Facebook, Google and IBM), although those capabilities have been developed for particular enterprise purposes. There's capability within the highly classified government world, also, which enables cyber and other security activities. Policies, strategies and principles have lagged AI development and application in the big-tech world—a phenomenon best captured by the Facebook 'move fast and break things' mantra.[16] Fortunately, that mindset hasn't applied to AI in the defence realm, where 'breaking things' is less forgivable, given that those things can be people. A key constraint on AI's application to national security has been the adversarial relationship between government and big tech in the US. This shows some signs of easing but not ending (for example, there's been a revival of anti-trust thinking).[17]

China is the other potential AI superpower through a combination of its state-centred data model and its own big-tech corporate sector.[18] China is also the widest state applier of data and tech for particular state purposes—including state surveillance of its population (think Xinjiang and China's social credit system) and state-centred data laws.[19] To the extent that all data is open to the state and the state can enable its tech-world actors to use it, China has a 'data advantage' that should translate into an AI advantage. However, Beijing's moves to reassert Chinese Communist Party control over big tech risk damaging it.[20]

But data is only part of an AI capability. Applying AI is a multidisciplinary team sport, and it turns out that datasets collected for particular purposes and in particular ways can have biases and limitations when used for other purposes. This is an issue for entities that have high risk appetites for rolling AI applications out, particularly for military or offensive cyber uses. Those who apply AI to weapon systems without deep understanding of the intended purpose, the environment and data limitations and a level of knowledgeable human participation are likely to inflict and experience nasty surprises.

Other states and supranational entities (such as the European Union) have capabilities, but not at the scale of the US or China. The US alliance system could enable states such as Australia to both contribute to and draw from US capabilities, while China's model is likely to remain a national one. Other nations and entities tend to be AI policy- and strategy-heavy,[21] with a large focus on getting ethical principles right,[22] but are short on applied capability that might use those policies and principles.

The national security implications of this for Australia are broad and complicated but, boiled down, mean one thing: if Australia doesn't partner with and contribute to the US as an AI superpower, it's likely to be a victim of the Chinese AI superpower and just an AI customer of the US. AUKUS is a step towards this AI partnership for national security.[23]

Part 4: Myths and mysteries

Is AI an existential threat to humanity?

Professor Anton van den Hengel, FTSE

Director, Centre for Augmented Reasoning, the University of Adelaide

There's a general and widely held misconception about AI. It's assumed that machines will be able to think and act like human beings, if not now, then soon. We're in fact a long, long way from that eventuality, if indeed we'll ever get to that point. And that's because AI poses a series of very difficult problems to solve.

The technology we have today is machine learning, which teaches a computer to do something specific, but only for one action or item at a time. While that performance, for example describing an image, can be amazing, it's still only a one-problem solution.

AI isn't what you think it might be and, besides some clever effects created in movies, it might never be the threat imagined by some. Here's why.

Take the example of a bee. A bee lives and functions in its environment every day—it collects food, is part of a community of bees and has a specialist role within the hive. Take that bee and put it in another environment, far removed from its home, and it will find food, create a nest and eventually find other bees with which to create a colony.

Your mobile phone has something like 100 bits of machine learning in it (and marketed as 'AI')—but leave it in a field far away from you and, after a few hours, the battery runs out. The device dies. It can't do anything by itself or to sustain itself.

AI, as the movies portray it, is far into the future, if we ever get to it at all. Yet the idea that AI can somehow turn conscious, and also somehow turn evil, is no doubt an exciting movie plot.

What such plots are based on is in reality a simple idea: that an AI's goals are misaligned with the goals of humanity. There's a famous example used to show how this works. Program a powerful AI-enabled machine to create paper clips in the most efficient manner possible. Without a limit on the number of paper clips needed, the machine will continue to build clips from existing materials and then mine resources and take over manufacturing until it has exhausted every possibility of manufacture from every possible mineral. In this theoretical story, the machine may well have exterminated all life forms that get in the way of making paper clips or could be used in the paper-clip manufacturing process.

The story shows that misaligned goals are the key to the dangers of AI. Machines have jobs to perform; in other words, they have goals. It isn't the machine that's the problem—it's the human making the decisions about the machine's goals that's the real issue. Just because AI may be 'smarter' in the way in which it solves problems at much faster speed than a human can, it doesn't follow that it controls anything outside of its own programming. The core of the issue, and the possibility of danger, is about humans making decisions about the goals intended for an AI platform.

This is most concerning when it comes to AI and warfare. There's a real fear that AI-enabled autonomous weapons systems will become the Skynet of *The Terminator* film franchise and that AI weapons systems will turn on their creators. Leaving aside the questions about the inability of AI to become self-aware, the real problem that faces us is about how and where such weapons are to be used. The real threat arises from the people controlling the mission goals of those systems and the parameters of action that AI autonomous weapons systems will have programmed into their software.

Yet there's a more immediate, real risk in AI and machine learning that's far more insidious. It has created a new form of power that's unevenly distributed. We've given away our data to essentially three companies that don't solve human problems. Those firms don't use or create AI and machine learning to solve the plague of domestic violence, or find ways to solve the carbon/climate crisis, or investigate ways to reduce Indigenous disadvantage and death. Rather, your data is churned through machine-learning algorithms to recommend to you the next product to buy.

In other words, the power of machine learning is about making money within monopolistic enterprises, which don't pay taxes to sustain the people and societies whose data they obtain and resell.

The real worry about AI isn't that *The Terminator* is ready to be built. That will probably never happen. It's a distraction from the main issue: that of the ethical use of AI as machine learning in mediating the very human problems that challenge us today.

Will machines ever replace people, and would that lead to mass joblessness?

Adam Slonim

Adjunct Fellow at Victoria University; Director of AI start-up Vestia.AI; consultant to Australian Institute for Machine Learning

Automation is the ability of machines to perform jobs that humans typically perform. Machine automation has been on the rise since 1139, when the Catholic Church tried (unsuccessfully) to ban the use of crossbows in war. Crossbows automated what was otherwise a specialist skill set by enabling untrained peasants to kill armoured knights on the battlefield with a very simple squeeze of a trigger.

AI takes automation a large step further. It makes predictions based on large and complicated amounts of data (which humans can do, although more slowly than computers) that are then used to create goods and services beyond human capabilities, such as searching the web for a story, picture or person, and to create whole new process chains and industries based on technologies we've never seen before, such as blockchain.

AI brings four primary effects to any economy:

- *The displacement effect*. New technologies can lead to a substitution of jobs and tasks currently performed by workers. Job displacement has the loudest volume in the debate about AI. And that's reasonable. The fear of lost jobs is real, and it is and will continue to occur, as it has occurred throughout history when technology overtook human endeavour, most specifically in the 18th century with the advent of the Industrial Revolution. In Australia, in particular, we still suffer nightmares from the vast displacement of jobs in the textile industry during the 1990s. Even though that displacement was caused by tariff reductions and the quest for low production costs—not by technological change—the fear of losing jobs is still high on our list of worries as a nation. Even putting this history into its appropriate historical context, the displacement effect is still not an absolute. Even with the closure of both the textile and car-manufacturing industries, Australia's GDP and standard of living have still risen. How come?
- *The skill complementarity effect*. There's a complementary increase in jobs and tasks necessary to use, supervise and grow from newly created technologies. At the turn of the 19th century, horses and buggies were *the* mode

of efficient and quick transport. When steam engines and cars took over, farriers and blacksmiths—once part of dominant industries—were now in tiny craft-based occupations. Now look at the jobs involved in producing and using cars and locomotives and project forward to driverless vehicles: the requisite number of jobs to support autonomous vehicle software, manufacturing and supply chains will be enormous. The classic case for the skill complementarity effect is the switch in Germany from coalmining to coal machine manufacturing, which in turn became a turbo-boost to the development of a machining-led export-oriented economy. When cars are made in Germany, they're certainly not made with low costs of labour—a lesson Australia never fully integrated.

- *The productivity effect*. Lower prices and higher disposable incomes drive increases in consumer demand. Look at the burgeoning number of cafes, fitness centres and speciality retail stores, which has ballooned in the past decade, and the rise in health-care, social-assistance, aged-care and infant-care jobs. Those sectors have experienced double-digit growth rates over the past decade.
- *The low cost of capital effect*. AI is increasingly embedded in newer and cheaper products accessible to more people than ever before. Have you used apps on your smartphone today? Got an idea for a small business and know it won't cost much to get an app rolled out? The steep fall in the cost of technology is driving increased access to technology and enlarging possibilities for innovation, in turn driving new forms of higher productivity. This is particularly beneficial for some economic sectors, such as agriculture, in which the application of technology has delivered enormous increases in crop yields at much lower costs of production.

Figure 5: New technologies can lead to a substitution of jobs and tasks currently performed by workers, so concerns about job displacement feature prominently in the debate about AI

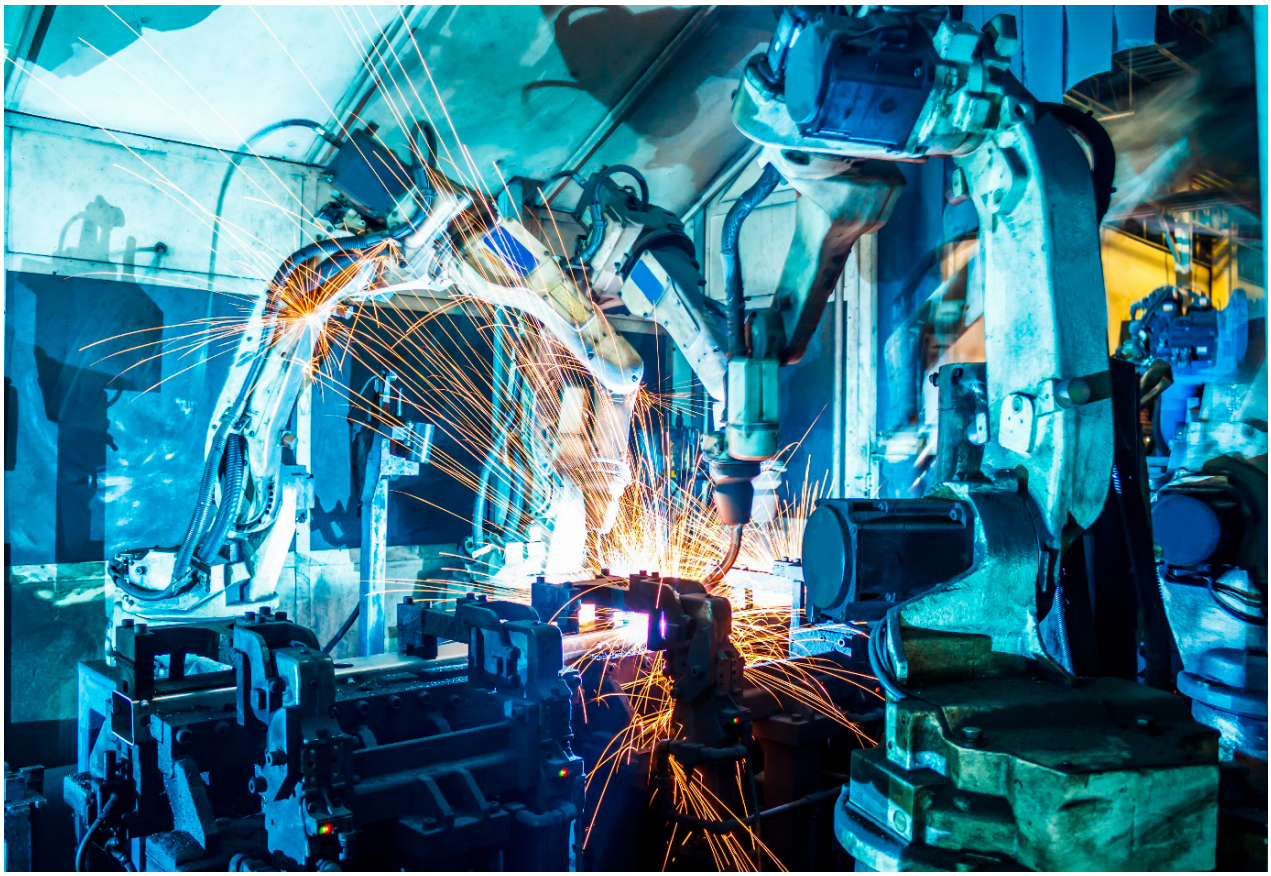

Source: iStockphoto, online.

What can Australia do to maximise the opportunities of AI and ward off threats of joblessness?

- Invest in technology for existing businesses, especially manufacturing.
- Ensure equal access to digital infrastructure, especially in a Covid-19-normal world in which access to online services, work and health is now crucial.
- Boost support for new AI businesses.
- Recapture old markets, especially for manufactured goods that new technology can produce more efficiently.
- Increase self-sufficiency in national security.
- Educate and reskill at all levels of society.
- Invest better, smarter and more heavily in AI R&D.

The workforce impact of AI is akin to the impact of computers and the internet. Most future jobs will have an AI component. Some jobs will disappear, and new jobs will be created. Eventually, it will be unimaginable to work without AI, much like it is today to work without a computer or smartphone. The challenge is that the jobs that will be created will be software- and service-oriented, and online. This means that those jobs will be located where the technology is, not where the customer is. This is a fundamental shift in how the economy is being structured in a global online world. Facebook, for example, employs more than 58,000 people but has only 140 employees in Australia.

It also means that foreign companies will increasingly use Australian data to extract income overseas. The question for Australia is whether we want to actively participate in this new global information-based and AI-driven market, or succumb to it.

How can we best manage ethical issues related to AI?

Dr Catriona Wallace

Chief Executive Officer, Ethical AI Advisory

The vast benefits of AI will be matched by its potential perils, especially in relation to health care, security and defence. The AI industry is largely without legislation and regulation, making it extremely dangerous. In fact, existential risk researcher Toby Ord puts the risk of AI being the most destructive force for humanity well above climate change and nuclear war, giving AI a 1-in-10 chance of destroying or fundamentally reducing humanity's potential by the end of this century.

What do we do about it?

We start by understanding AI ethics and extending that understanding to the existential risks that AI poses.

Much work has been done in the field of ethical AI and, more recently, responsible AI.[24] The Australian Responsible AI Index report (n = 416) released by Ethical AI Advisory and Gradient Institute shows that Australian-based AI organisations are at an 'initiating' level of maturity. That equates to about a 60/100 performance score and suggests that Australian-based organisations have done some planning for AI and are initiating some AI projects. The proportion of Australian organisations that were classified as 'mature' was 8%. The leading industries for maturity in responsible AI were manufacturing and finance and services. The least mature industries were construction and retail and hospitality. Only about 30% of participants in the study had actioned any of the plans relating to responsible AI, and only 57% were aware of the government's ethical AI principles.

Hence, Australia has much to do to close the responsible AI gap.

The best way to manage ethical AI problems is to first reduce the chance that they'll occur. To do that, organisations should follow principles that include the following:

- AI should be built with humans, society and the environment in mind.
- AI should be built with human-centred values at its core.
- AI should be fair and not discriminate.
- AI should be reliable and safe.
- AI should adhere to security and privacy requirements.
- AI should be contestable.
- AI should be transparent and explainable.
- Those developing and deploying AI should be held accountable for any harm caused.

To get started on the responsible and ethical AI journey, organisations should:

1. *Conduct stakeholder engagement*

   Consult widely with stakeholders to identify the harms the system may cause, then determine ethical objectives to control those harms and ensure the AI system achieves them.

2. *Identify risks*

   Identify the people at risk of being systematically disadvantaged by the AI system and ensure that special consideration is given to protecting them.

3. *Document AI systems*

   Document AI systems that can affect the lives of people—their purpose, risks, key design decisions and justifications, and performance, and who is responsible for them.

4. *Continuously monitor AI systems*

   Continuously monitor AI systems against their business and ethical objectives, search for unintended harms, and build in mechanisms for review, redress and mitigation.

5. *Train staff*

   Train staff in the novel risks of AI systems and their roles in controlling those risks.

6. *Manage risks*

   Extend existing risk-management frameworks to incorporate risks introduced or potentially amplified using AI systems.

If we as a nation were to really double down on responsible AI, Australia could develop itself as a trusted AI nation, which will be particularly important as AI drives further societal innovation, such as the metaverse.

Do we need social scientists and the general public to contribute to AI development?

Professor Anna Ma-Wyatt

Professor, School of Psychology, the University of Adelaide; Co-Director, CNRS IRL CROSSING

To most people, social science and computer science seem like they're at opposite ends of a spectrum. But social science—especially behavioural science—is really about ways of understanding how humans interact with each other and the world around them. There are several ways in which social scientists and the general public can contribute to AI development and, if we get this right, I think it would be of great benefit.

Let's start with thinking about how social science might contribute to AI development. In behavioural science, we've traditionally used lab-based studies and very careful measurements of behaviour. Those approaches have allowed us to build theories about how mechanisms for all kinds of human behaviour work—from memory, to visual processing, to language.

With advances in sensors and sensing, there are new and exciting ways to measure human behaviour. AI and machine-learning techniques are already commonly used to help process high-volume datasets, and that approach has been especially effective in behavioural neuroscience.

Moreover, advances in sensors and sensing also offer new opportunities to measure human performance outside of the lab and as people interact with their environment in more naturalistic and more complex ways. These types of approaches produce enormous amounts of data, so AI can be helpful for analysing that data and helping us work out new ways of understanding behaviour. They've already shown their power by allowing us to understand patterns of movement of older people at home.

As AI becomes more prevalent in everyday life, understanding how humans interact with automated systems is another exciting multidisciplinary area. For example, cutting-edge approaches to cybersecurity acknowledge that human behaviour can be a huge risk, so aspects of behavioural science are being integrated into those approaches. Teams consisting of social scientists, political scientists, lawyers, data scientists and computer scientists are already working together.

A key part of the success of this approach is making sure we have good dialogue between behavioural and social scientists and people developing AI. We need a respectful environment in which we can explore how to integrate our approaches to have greater explanatory and predictive power. If social science and AI can work together, they offer exciting opportunities for developing new ways of understanding humans and developing new forms of technology that humans can use in different contexts.

The broader community has a truly important role to play not only as users and consumers, but also because we as a society need to understand the implications of AI. We need to make sure developers help the general public (who are users, too) understand the potential of AI but also the implications of using it for different purposes. There can be a lot of confusion around any new technology. To counter that, it can be helpful for the public to be engaged in understanding the new opportunities and understanding the strengths and limitations of AI and how it can be applied. The education of users and the ability to communicate clearly with the public will be important skills for people working in AI development.

There are also significant questions about the ethical use of AI that would benefit from having the general public and social scientists, among others, as part of those discussions. AI will be so pervasive that our conversations must be inclusive and broad across all levels of society to ensure that we use an ethical approach that fits with our culture. Public policy and law will of course have a very important role to play in those conversations.

We need to make sure that there's access across all parts of society, but how do we ensure that? What do 'access' and 'representation' mean? What are the sociotechnical and infrastructure requirements necessary to make this possible? Social scientists (for example, anthropologists, psychologists, sociologists) will be invaluable in working with teams to understand the social implications and working with AI developers and policy- and lawmakers. For example, what sort of cultural norms are OK to have in AI? Who in society gets to make those decisions? There are already great groups working on those questions (such as the 3A Institute at the Australian National University), and I think it's important work that needs to be part of the conversation.

Part 5: The future

Why must Australia invest in AI research and development?

Misha Schubert
Chief Executive Officer, Science & Technology Australia

Most technological advances spark incremental progress. And then a few come along that are truly game-changing.

AI is one of those truly transformative technologies. It's set to revolutionise our lives and workplaces at rapid speed in the coming decade. It will potentially reshape almost every job, industry and life. It's already being deployed across a wide array of Australian industries and sectors.

In the legal profession, AI is now being used to scan and prepare legal briefs.[25] There's also a growing array of analytics applications being used to save people time and resources.[26]

In medicine and health care, the potential of AI is enormous, but yet to be truly realised.[27] Many parts of health care stand to benefit from a data-driven approach, including image analysis to improve the detection and diagnosis of disease.[28] Analysis of patient data has helped to predict and prevent falls among at-risk patients.[29] However, these game-changing possibilities need more work and more investment to reach their full lifesaving potential.

A uniquely Australian example can be found in Kakadu National Park. Indigenous owners are working with researchers to support ecosystem resilience through a combination of Indigenous knowledge, technology and AI.[30] Under the direction of Indigenous rangers, drone footage collects data over areas of the national park that are difficult to access. AI is used to scan hours of video footage to identify para grass—an introduced fodder crop that's displacing the native grasses, which are a key habitat for magpie geese. Strong numbers of magpie geese indicate that the country is healthy. Adaptive management has already seen the population in one wetland increase from 50 to 1,800 birds.[31]

So why should Australia invest in AI R&D? If we want to secure the economic opportunities that will flow from AI, and shape those frontier technologies as they evolve, we must invest in a specialised, skilled workforce and advances in key technologies.[32] If we don't, we'll be left behind—and our children and grandchildren will miss out on vast opportunities.

Australia's *Artificial intelligence roadmap* cites three pivotal areas in which AI has the potential to shape our nation: natural resources and the environment; health, aging and disability; and cities, towns and infrastructure.[33] Those diverse fields highlight the potential breadth of AI applications in widely divergent aspects of our lives.

Other countries are investing, and investing big.

The biggest players in the AI game are the US and China.

The American National AI Initiative was unveiled in 2019.[34] It drives the safe development, testing and deployment of AI technologies. The strategy recognises the importance of international collaboration, while striving to put the US at the front of global AI research. The US Government's investment in AI projects amounted to around US$5 billion in 2020.[35] In 2021, the National Artificial Intelligence Initiative Office was launched to play a coordinating role for all AI research and policymaking across academia, government and industry.[36]

China's strategy, announced in 2017, aims to make China an AI leader by 2030. It has a key focus on technologies such as unpiloted aerial vehicles and voice and image recognition. While exact government funding amounts aren't known, it's clear that China is significantly upping its game in research activity, with very steep growth in publications on AI.[37]

The Canadian Government announced a national AI plan in 2017, investing C$125 million (~A$137 million) to support the Pan-Canadian Artificial Intelligence Strategy.[38] The strategy sets ambitious goals to attract and retain world-class AI researchers, nurture a collaborative AI ecosystem, support national AI initiatives, and contextualise AI work in its ethical and societal implications.

In 2018, the UK pumped nearly £1 billion (~A$1.87 billion) into AI research.[39] It released its National AI Strategy in September 2021—a '10-year plan to make Britain a global AI superpower'.

In South Korea, the government allocated ₩1.7 trillion (~A$1.1 billion) in 2020 for data, networks and AI under its National Strategy for Artificial Intelligence.[40] The strategy includes nine plans to drive advancements in the AI ecosystem, AI use and people-centred AI.

Singapore released its National AI Strategy in 2019.[41] Its goal is to propel 'Singapore as a leader in developing and deploying scalable, impactful AI solutions, in key sectors of high value and relevance to our citizens and businesses.' The Singaporean Government has supported the strategy with funding totalling S$680 million (~A$695 million). The money will support projects in fundamental AI research, translational research, and industry–research collaborations.

India is a very strong player, too. Its government released an AI strategy in 2018, identifying areas of focus in agriculture, health and education.[42] This was backed by ₹7,000 crore (A$1.3 billion). The Indian Government is also working on its National Programme for Artificial Intelligence, with a goal of 'making India the global leader in AI, ensuring responsible and transformational AI for all.'[43]

In the Australian context, the potential uses of AI to create safer and more interesting jobs and take on tasks that are either dangerous or repetitious are almost limitless. And the technology itself is even evolving our legal concepts.

In a landmark judgement in August 2021, Australia's High Court ruled that an AI can be named as an inventor in a patent application—a decision that challenges the prevailing thought that AI systems can't be truly creative.[44]

As the world strives to deepen AI capabilities, diversity must be a critical policy consideration. AI has the power to radically transform our society, and we need to think constantly and simultaneously about the ethics and social implications as we develop the research. We need the best minds, of all genders, focused on this task. At the moment, only 22% of AI researchers are women.[45] Given the huge ethical and societal implications of AI research and its applications, we need to ensure a diversity of brainpower, perspectives and experience to ensure optimal outcomes for our societies.

Diversity of thought also comes from including and incorporating other world views and cultural perspectives into AI research. Indigenous knowledge systems contain a wealth of experience, particularly to understand and live sustainably within our environment. Including Indigenous knowledge to inform AI research will lead to better, more holistic and more inclusive AI.[46]

The CSIRO and Data61 forecast that AI benefits will be worth $22.17 trillion to the global economy by 2030.[47] They predict that Australia could boost its economy by $315 billion by 2028 from digital technologies including AI.

The CSIRO says that, in the next decade, by deepening our use of digital technologies, we could snare a slice of $30–50 billion in future Asia–Pacific markets for healthcare innovations, prevent up to 1,100 road deaths a year, and boost the farm-gate value of agriculture by $20 billion.

It will be a race to secure the benefits AI has to offer societies and nations. Countries around the globe are setting forward-looking strategies, with big ambitions, and backing them with significant funding. Australia needs to be in the race and to invest—at scale, and ensuring broad inclusion—in AI R&D. It's a powerful investment in securing our future prosperity.

How do we ensure that our workforce is ready for AI?

Senator Rex Patrick

Senator for South Australia

AI offers the Australian workforce huge opportunities—if we do it right. 'Right' means implementing AI that sits alongside us to deliver workplace efficiencies and free up human ingenuity for tasks that technology can't tackle. It means democratising AI and co-designing and delivering innovation for people across all parts of society. It means creating sovereign capability and building our own solutions for our own problems—not purchasing and retrofitting technology that other countries have developed. 'Right' means building complexity and resilience into our economy.

It's important to acknowledge the fears some members of our community have about AI in the workplace. The key here is to recognise that we can manage how and what we adopt by being proactive about our own AI agenda. AI is already operational in our phones, our smart technologies, in our cars, in our buildings. Its presence in our working lives will increase whether we like it or not, so it's in our national interest to take control of Australia's AI.

Creating a workforce that's ready for AI involves starting from the ground up at all levels of industry and education.

First, we must transition our businesses away from being AI consumers towards being AI innovators. The Australian Government has several mechanisms at its disposal to create incentives and stimuli to encourage business to innovate with AI. Business itself will need to come to grips with AI as a core part of its processes and the delivery of high-quality customer experiences.

Second, we should invest in bridging the gap between Australia's considerable academic AI capability and the needs of our industries. There are two parts to this problem: we don't train enough AI-capable people in Australia; and those we do train are attracted to go overseas, where more interesting and high-paying jobs exist.

The solution to this problem is to fund more places at Australian universities and to create more flexible working arrangements for AI-trained personnel to work in commercial enterprises, government agencies, or both, while still at university. It's here that we can look to the US and Israel and replicate the right elements of these flexibilities to ensure that we create our own innovative AI solutions and train our own AI talent.

Third, we can create bespoke AI solutions and tools for Australia's resources, agriculture, education and tourism sectors. Augmented reasoning is an exciting new field of AI that offers exciting potential for our workplaces through its elevated capacity for understanding human instructions and needs, and for achieving greater social acceptance.

More funding through existing research channels, such as Australian Research Council grants and R&D tax incentives, will help, but the emphasis has to be on leading-edge technologies that give Australia competitive and comparative advantages.

Fourth, we need to study how to achieve and then implement the learning and training of AI within our TAFE sector and even in our secondary education systems. There are several practical reasons for doing this. AI is becoming part of our everyday existence and is already embedded in the technologies we use daily.

Figure 6: AI offers Australia's future workforce huge opportunities, if we do it right

Source: iStockphoto, online.

AI applications are coming that will allow builders, for example, to use smart efficiencies in their processes and for waste companies to turn a more diverse range of waste products into renewable-energy-based systems. Builders will have more and more componentry quality-assessed by AI applications, and more materials to be used more efficiently. Plumbers will have AI-enabled tools to assess features and defects of water and gas lines. Maintenance providers will use AI to predict building and system requirements—not just wait for things to go wrong—saving enormous amounts of money and time and, in so doing, extending the life cycle of materials. Supply chains and warehouses are being radically reshaped by AI logistics systems.

Every industry will be affected by AI, and every job will have an AI component to it. AI will increasingly become a foundation of our economy and society. Do we leave this opportunity for some other country or overseas tech firm to build for Australia, or take up the mantle for ourselves?

The jobs of the future will incorporate AI, not be replaced by it. We must take deliberate actions to get there and, in so doing, become a maker of technology that both enhances our sovereign security and becomes a world-class source of export opportunity.

How does government enable and/or use industry as a multiplier for AI projects?

Adam Reid
Chief Executive Officer, Department for Innovation and Skills, Government of South Australia

Traditional models of innovation are typically based on new ideas and inventions being exploited by new businesses that can disrupt markets and gain competitive advantage over incumbent enterprises. We've seen many examples of businesses, including technology businesses, that have lost their once-dominant position due to emerging technologies and competition from more innovative and nimble businesses that are taking risks and investing in disruptive technologies to gain market share. The challenge remains to convince established businesses to realise the value of investing in innovation or risk losing their competitive advantage to more agile enterprises with a greater appetite for risk, and reward.

Australian business investment in R&D and innovation is largely heading in the wrong direction. Business expenditure on R&D sits at under 1% as a share of GDP, down from a historical high of 1.37% in 2008–09. Some of the world's most advanced economies are recording business expenditure on R&D of over 2%, nearing 4% of GDP in the case of Israel.

Although the share of innovation-active businesses has risen in the past decade in Australia, relatively few of those enterprises are producing innovations that are new to Australia and the world, and our share of business expenditure on innovation mirrors our performance on R&D compared with the rest of the world. The performance of South Australia's business sector is broadly consistent with those national trends.

Clearly, a continuation of this situation is going to have a negative impact on our current high standards of living as competitive advantage and the rewards from investing in innovation shift to our competitors overseas. The challenge facing Australian governments is obvious, although the solutions are often complex—especially given the significant shift in culture required of industry and often government itself.

In the past few years, the South Australian Government has been working to encourage greater business engagement in R&D and innovation by facilitating collaborative partnerships between the research sector and industry. Our 'EXCITE' Science Strategy aims to connect ideas and expertise with small and large businesses through initiatives such as the Innovation Challenge on Augmenting Ability in partnership with MTP Connect, which will bring together businesses, researchers and innovators to develop cutting-edge products in the health, ageing and disability sectors.

Our external innovation and translation intermediaries will work to bridge the gaps in information and access that can hamper innovation and will drive collaboration between businesses and researchers within the state's innovation districts. The government has committed up to $3 million to appoint our first intermediary organisation for Adelaide BioMed City, and further intermediaries are planned for our innovation districts at Tonsley in Adelaide's south and Lot Fourteen in central Adelaide.

EXCITE's Frontier Technology Centres program will work with small and medium-sized enterprises (SMEs) and business leaders to harness the capability of emerging technologies. As a precursor to that program, the state government made strategic investments to grow leading-edge capabilities in key strategic technologies, such as photonics and advanced sensing, satellite technologies, cybersecurity, and AI.

While AI has been part of our daily lives since at least the 1980s, the surge in awareness—driven by the growing digitalisation of business and processes, increases in computing power and access to large datasets for training AI algorithms—is driving major new investment in AI technologies, including in South Australia.

The state has seen the benefits of long-term investment in AI, stretching back to the government's initial investment in the Australian Centre for Visual Technologies (ACVT) as early as 2007, in collaboration with the (then) Defence Science and Technology Organisation and other industry partners. That relatively modest investment provided the impetus for ACVT to forge partnerships beyond defence and security and to work with a range of diverse businesses, such as the Victorian sports statistics company Champion Data, medical image analysis company LBT Innovations, simulation company Sydac (now part of the global Sogeclair Group) and digital media company Monkeystack.

This approach to industry engagement by ACVT continued to grow and evolve through the Australian Institute for Machine Learning (AIML), which launched in 2018 with a significant investment of $7.1 million from the state government. Based at Adelaide's latest innovation district, Lot Fourteen, AIML has committed to renewed engagement with industry and government, which has seen the institute grow in size and capability to become the largest university-based machine-learning research group in Australia and one of the global top five sites for computer vision research output.

The calibre of AIML's research has attracted the attention of some of the world's largest multinational corporations, such as Amazon, Accenture, Deloitte, Microsoft Azure, MTX, Google Cloud and Nokia 5G, all of which have established a presence in Adelaide to collaborate with researchers at AIML.

Importantly, the government's investment in AIML is supporting the institute to engage with SMEs that are looking to automate and harness AI to gain competitive advantage:

- Adelaide visual-effects studio Rising Sun Pictures has produced special effects for some of Hollywood's biggest blockbuster films and worked with AIML to develop novel VFX tools that have enabled Rising Sun to deliver timely and superior results to global production studios.
- Geospatial tech company Aerometrex used AIML's deep-learning capabilities to develop highly accurate 3D maps that can reveal shadows cast by buildings at different times of the day, helping planners and developers to gauge the impact on surrounding areas.
- Regulatory technology business Neo Analytics worked with AIML to apply machine-learning models to improve its regulatory compliance monitoring software to deal with significant amounts of data necessary for financial institutions to carry out their business.

As well as supporting SMEs to explore the benefits of AI, government has a role to play as a regulator and service provider by exploring uses of AI to improve accuracy and efficiency across a range of functions.

The state's Department of Primary Industries and Regions worked with AIML to apply machine learning to satellite imagery to assist land condition monitoring across South Australian pastoral leases. AI is also being applied in the resources sector as the state's Department for Energy and Mining works with AIML researchers to apply deep learning to South Australia's vast collection of geological survey data. The work has resulted in a model to predict rock outcrops across the state, which will benefit mineral exploration and help South Australia retain its position as a top 10 destination in the world for mining investment attractiveness.

Clearly, the government's early investment in frontier AI technology is starting to bear fruit as the state sees a nascent community of more than 50 AI and AI-enabled businesses, ranging from start-ups to global multinationals, tapping into world-class research and extensive datasets and applying analytical tools to get the edge in competitive markets.

Our ongoing challenge is to convey the benefits of innovation, including the potential benefits of AI and automation, to the broader business sector. By demystifying the technology and facilitating networks across business, academia and government, South Australia is building a strong ecosystem that will drive employment and exports and attract talent to work on the next generation of disruptive technologies.

Notes

1 Prime Minister, Minister for Defence, Minister for Foreign Affairs, Minister for Women, 'Australia to pursue nuclear-powered submarines through new trilateral enhanced security partnership', media statement, 16 September 2021, online.
2 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, Li Fei-Fei, 'ImageNet: A large-scale hierarchical image database', IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2009, online.
3 H Lv, L Shi, J Berkenpas, F Dao, H Zulfiqar, H Ding, Y Zhang, L Yang, R Cao, 'Application of artificial intelligence and machine learning for COVID-19 drug discovery and vaccine design', *Briefings in Bioinformatics*, 2021, 22(6).
4 Organisation for Economic Development and Co-operation, 'Gross domestic spending on R&D', online.
5 *Harvard Business Review*, December 2018.
6 'Worldwide spending on artificial intelligence is expected to double in four years, reaching $110 billion in 2024, according to new IDC spending guide', IDC Corporate USA, 25 August 2020, online.
7 C-suite = chief executive officers, chief financial officers, chief information officers and so on.
8 SurvivAI, 2021.
9 Critical Technologies Policy Coordination Office, *The Action Plan for Critical Technologies*, Department of the Prime Minister and Cabinet, Australian Government, 2021, 15–16, online.
10 Department of Defence, *More together: Defence Science and Technology Strategy 2030,* Australian Government, 2019, online.
11 SA Hajkowicz, S Karimi, T Wark, C Chen, M Evans, N Rens, D Dawson, A Charlton, T Brennan, C Moffatt, S Srikumar, KJ Tong, *Artificial intelligence: solving problems, growing the economy and improving our quality of life*, CSIRO Data61, Commonwealth Scientific and Industrial Research Organisation, Australia, 2019, online.
12 Daniel Zhang, Saurabh Mishra, Erik Brynjolfsson, John Etchemendy, Deep Ganguli, Barbara Grosz, Terah Lyons, James Manyika, Juan Carlos Niebles, Michael Sellitto, Yoav Shoham, Jack Clark, Raymond Perrault, *The AI Index 2021 annual report*, AI Index Steering Committee, Human-Centered AI Institute, Stanford University, Stanford, California, March 2021, online.
13 Department of Defence, *Defence Data Strategy 2021–2023*, Australian Government, 2021, 35, online.
14 Hajkowicz et al., *Artificial intelligence: solving problems, growing the economy and improving our quality of life*.
15 Standards Australia, *An artificial intelligence standards roadmap: making Australia's voice heard: final report*, 2020, online.
16 Hemant Taneja, 'The era of "move fast and break things" is over', *Harvard Business Review*, 22 January 2019, online.
17 Nicolas Rivero, 'A cheat sheet to all of the antitrust cases against Big Tech in 2021', *Quartz*, 29 September 2021, online.
18 *Mapping China's tech giants*, ASPI, Canberra, 2022, online.
19 Katja Drinhausen, Vincent Brussee, *China's social credit system in 2021: from fragmentation towards integration*, MERICS: Mercator Institute for China Studies, 3 March 2021, online.
20 Lulu Yilun Chen, Jun Luo, Zheng Li, 'China crushed Jack Ma, and his fintech rivals are next', *Bloomberg*, 24 June 2021, online.
21 'A European approach to artificial intelligence', European Commission, no date, online.
22 Denham Sadler, 'Alan Finkel on AI ethics and law', *InnovationAus.com*, 10 December 2019, online.
23 'Joint Leaders Statement on AUKUS', The White House, 15 September 2021, online.
24 Australian Human Rights Commission, *Human rights and technology: final report*, September 2021, online.
25 'Legal research, minus the lengthy search', *Judicata*, no date, online.
26 Daniel Faggella, *AI in law and legal practice—a comprehensive view of 35 current applications*, Emerj Artificial Intelligence Research, 7 September 2021, online.
27 Bertalan Mesko, Marton Gorog, 'A short guide for medical professionals in the era of artificial intelligence', *npj Digital Medicine*, 24 September 2020, 3(126), online.
28 Daniel Greenfield, 'Artificial intelligence in medicine: applications, implications, and limitations', *SITN: Science in the News*, Graduate School of Arts and Sciences, Harvard University, 19 June 2019, online.
29 Juliet Van Wagenen, 'Healthcare analytics point providers to patients that need the most care', *HealthTech*, 17 April 2017, online.
30 'AI transforms Kakadu management', news release, Microsoft, 20 November 2019, online.
31 Larry Marshall, 'AI + Indigenous knowledge a powerful tool posing critical questions', *CSIRO Algorithm*, 27 November 2019, online.

32 Department of Industry, Science, Energy and Resources, 'Techtonic 2.0—How to AI-proof our workforce', Australian Government, *YouTube*, no date, online.
33 Data61, *Artificial intelligence roadmap*, CSIRO, online.
34 National Artificial Intelligence Initiative, US Government, online.
35 Neil Savage, 'The race to the top among the world's leaders in artificial intelligence', *Nature*, 9 December 2020, online.
36 'The White House Launches the National Artificial Intelligence Initiative Office', news release, The White House, 12 January 2021, online.
37 Savage, 'The race to the top among the world's leaders in artificial intelligence'.
38 'Pan-Canadian AI Strategy', CIFAR, no date, online.
39 'National AI Strategy', UK Government, September 2021, online.
40 'National Strategy for Artificial Intelligence', South Korean Government, 28 October 2019, online.
41 'National Artificial Intelligence Strategy: Advancing our Smart Nation', Singaporean Government, November 2019, online.
42 'AI policy and national strategies', in *Artificial Intelligence Index report 2021*, Stanford University, 2021, online.
43 'India's AI marathon has begun well', *Hindustan Times*, 18 September 2021, online.
44 Alexandra Jones, 'Artificial intelligence can now be recognised as an inventor after historic Australian court decision', *ABC News*, 1 August 2021, online.
45 Natalie Marchant, 'Only 22% women in AI jobs—the gender gap in science and technology, in numbers', *The Print*, 17 July 2021, online.
46 Rachael Bolton, 'Indigenous knowledge systems in the age of AI', *InnovationAus.com*, 3 December 2021, online.
47 Data61, *Artificial intelligence roadmap*.

Acronyms and abbreviations

ACVT	Australian Centre for Visual Technologies
AI	artificial intelligence
AIML	Australian Institute for Machine Learning
ASPI	Australian Strategic Policy Institute
DAIRNet	Defence Artificial Intelligence Research Network
GDP	gross domestic product
NASA	National Aeronautics and Space Administration
OECD	Organisation for Economic Co-operation and Development
R&D	research and development
SMEs	small and medium-sized enterprises
STEM	science, technology, engineering and mathematics
TAFE	technical and further education
VET	vocational education and training

Some recent ASPI publications

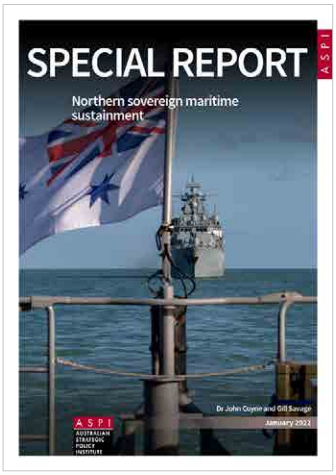

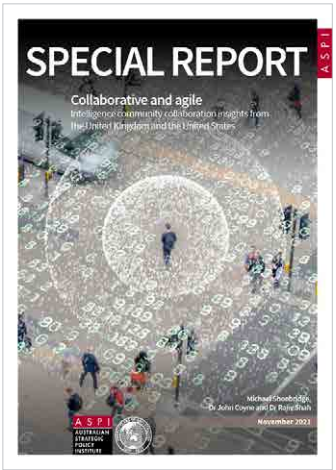

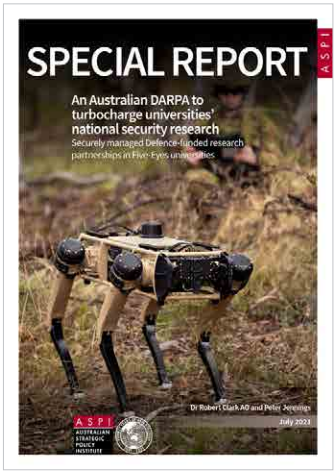

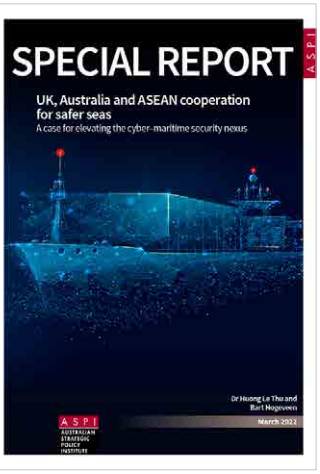

WHAT'S YOUR STRATEGY?

Stay informed via the field's leading think tank, the Australian Strategic Policy Institute.

The Strategist, ASPI's commentary and analysis website, delivers fresh ideas on Australia's defence and strategic policy choices as well as encouraging discussion and debate among interested stakeholders in the online strategy community. Visit and subscribe to an email digest at www.aspistrategist.org.au.

facebook.com/ASPI.org

@ASPI_org

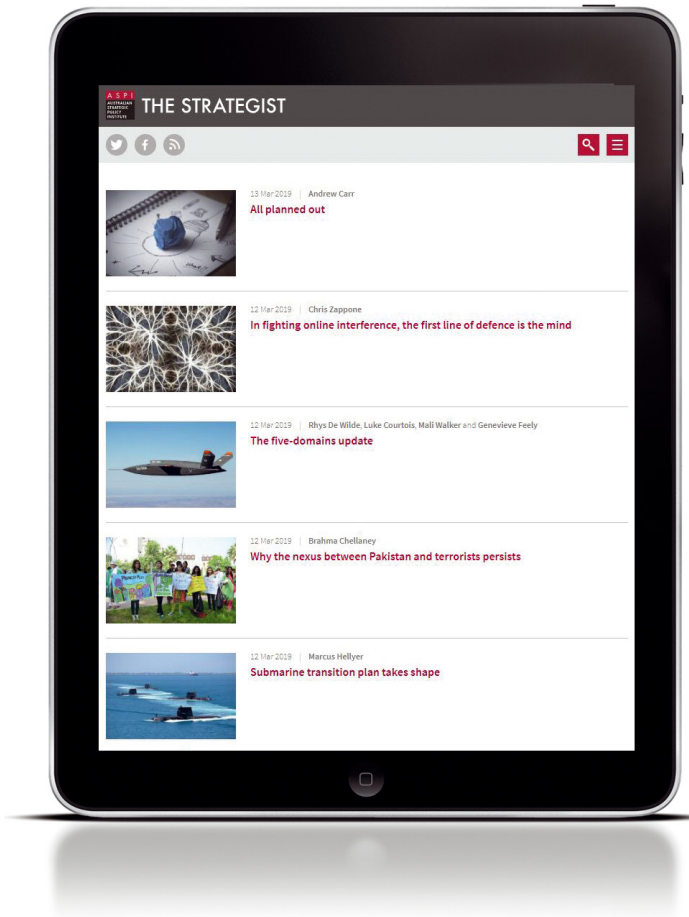

ASPI
AUSTRALIAN STRATEGIC POLICY INSTITUTE

Supported by

LOCKHEED MARTIN THALES

To find out more about ASPI go to www.aspi.org.au or contact us on 02 6270 5100 and enquiries@aspi.org.au.

Artificial intelligence
Your questions answered